

Best of TechEd 2012

Windows IT Pro

A PENTON PUBLICATION

AUGUST 2012 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Navigate VMware Licensing

What Would Microsoft Support Do?

Troubleshooting
Windows Server 2008 R2
Failover Clusters

Shared-Nothing Live Migration
with Windows Server 2012
Hyper-V

SharePoint Security 101

Use PowerShell for
Administrative Reporting

Understand System Center
App Controller 2012

Plus >>

vSphere

ONLY 1&1 OFFERS YOU THE RELIABILITY OF DUAL HOSTING



What is Dual Hosting?

Your website hosted across multiple servers in 2 different data centers, in 2 different geographic locations.

Dual Hosting,
maximum reliability.

1&1 – get more for your website!

- ✓ **More Possibilities:**
65 Click & Build applications.
- ✓ **More Included:**
Free domain*, **free** e-mail accounts, unlimited traffic,
NEW: Adobe® Dreamweaver® CS5.5* and much more.
- ✓ **More Privacy:**
Free private domain registration.
- ✓ **More Reliability:**
Maximum reliability through hosting simultaneously across two separate data centers.



SALE ENDS
AUGUST 31, 2012

ALL 1&1 HOSTING PACKAGES

\$3.99 per month
SAVE UP TO 60%!*

DOMAIN OFFERS: .COM/.ORG JUST \$ 3.99 (first year)*



www.1and1.com



* 12-month minimum contract term and 3-month pre-paid billing cycle apply for web hosting offer. Standard prices apply after first year for domain and hosting offers. Free domain with Unlimited and Business hosting packages. Dreamweaver included only in Business hosting package. Visit www.1and1.com for billing information and full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet, all other trademarks are the property of their respective owners. © 2012 1&1 Internet. All rights reserved.



**CREATING INTERNET FACING
SHAREPOINT SITES WITH
SHAREPOINT 2012**

SEPTEMBER 6TH

FREE VIRTUAL CONFERENCE

**SPEAKERS INCLUDE
TODD BAGINSKI, CELINA BAGINSKI
AND MORE!**

**CLICK ON THIS AD TO REGISTER
AND FOR MORE INFORMATION!**

COVER STORY ▼

Navigate VMware Licensing — Alan Sugano

52

Learn how best to manage the licensing changes in the most recent version of vSphere.

Features

60 **Shared-Nothing VM Live Migration with Windows Server 2012 Hyper-V**
John Savill

72 **Administrative Reporting with PowerShell**
Max Trinidad

85 **Understanding App Controller 2012**
Damir Dizdarevic

95 **SharePoint Security 101**
Randy Williams

Interact

44 **Ask the Experts**

In Every Issue

13 **IT Community Forum**

138 **Ctrl+Alt+Del**

139 **Advertiser Directory**

139 **Directory of Services**

139 **Vendor Directory**

Chat with Us



Facebook



Twitter



LinkedIn

Columns

7

IT Pro Perspectives

Windows Server 2012 Is Good News for IT

Sean Deuby



16

Need to Know

Windows and Visual Studio RCs and Office 15 Beta

Paul Thurrott



22

Windows Power Tools

Pipeline Problems and Get-ADUser

Mark Minasi



25

Top 10

Windows Server 2012 Storage Enhancements

Michael Otey



31

Enterprise Identity

Virtualization-Safe Active Directory in Windows Server 2012

Sean Deuby



36

What Would Microsoft Support Do?

Troubleshooting Windows Server 2008 R2 Failover Clusters

John Marlin



Products

105 New & Improved

109 Paul's Picks

Paul Thurrott

110 Best of TechEd 2012 Winners

Jason Bovberg

116 ControlPoint 4.5

Russell Smith

121 Avance

Joel Sloss

129 Apple iPad for the IT Pro

Michael Dragone

134 Industry Bytes

Editorial

Editorial Director:

Megan Keller

Editor in Chief:

Amy Eisenberg

Senior Technical Director:

Michael Otey

Technical Director:

Sean Deuby

Senior Technical Analyst:

Paul Thurrott

Custom Group Editorial Director:

Dave Bernard

Exchange & Outlook:

Brian Winstead

Systems Management,

Networking, Hardware:

Jason Bovberg

Scripting:

Blair Greenwood

Security, Virtualization:

Amy Eisenberg

SharePoint, Active Directory:

Caroline Marwitz

SQL Server, Developer Content:

Megan Keller

Managing Editor:

Lavon Peters

Assistant Managing Editor:

Rachel Koon

Editorial SEO Specialist:

Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,

Tony Redmond, Paul Robichaux,

Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,

Jeff Felling, Brett Hill, Dan Holme,

Darren Mar-Elia, Eric B. Rux,

William Sheldon, Curt Spanburgh,

Bill Stewart, Orin Thomas,

Douglas Toombs, Ethan Wilansky

Art & Production

Production Director: Linda Kirchgesler

Senior Graphic Designer: Matt Wiebe

Advertising Sales

Publisher: Peg Miller

Key Account Director:

Chrissy Ferraro • 970-203-2883

Account Executives:

Barbara Ritter • 858-367-8058

Cass Schulz • 858-357-7649

Client Project Managers

Michelle Andrews • 970-613-4964

Kim Eck • 970-203-2953

Ad Production Supervisor:

Glenda Vaught

Marketing & Circulation

Customer Service

Senior Director, Marketing Analytics:

Tricia Syed

Online Sales Development Director:

Amanda Phillips • 970-203-2806

Technology Group

Senior VP, Penton Media Technology Group:

Kim Paulsen

Corporate

Chief Executive Officer:

David Kieselstein

Chief Financial Officer/Executive Vice

President: Nicola Allais



List Rentals

MeritDirect

333 Westchester Avenue,

White Plains, NY 10604

Reprints

Reprint Sales:

Wright's Media • 877-652-5295

Windows IT Pro, August 2012, Issue no. 216, ISSN 1552-3136. *Windows IT Pro* is published monthly by Penton Media, Inc. Copyright ©2012 Penton Media, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any way without the written consent of Penton Media, Inc.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525, 800-621-1544 or 970-663-4700. Customer Service: 800-793-5697.

We welcome your comments and suggestions about the content of *Windows IT Pro*. We reserve the right to edit all submissions. Letters should include your name and address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming code in this issue is ©2012, Penton Media, Inc., all rights reserved. These programs may not be reproduced or distributed in any form without permission in writing from the publisher. It is the reader's responsibility to ensure procedures and techniques used from this publication are accurate and appropriate for the user's installation. No warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server® are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries and are used by Penton Media, Inc., under license from owner. *Windows IT Pro* is an independent publication not affiliated with Microsoft Corporation. Microsoft Corporation is not responsible in any way for the editorial policy or other contents of the publication.

Windows IT Pro

WINCONNECTIONS: OCT 29 - NOV 1 • BELLAGIO • LAS VEGAS

CLOUD
ONLINE

WINDOWS
ONLINE

Microsoft®
Exchange
ONLINE

SQL Server
ONLINE

SharePoint
ONLINE

KEYNOTES



PAUL THURROTT
WINDOWS IT PRO
Senior Technical Analyst



JEFFREY SNOVER
MICROSOFT
Distinguished Engineer and the Lead Architect for Windows Server



MARK MINASI
MINASI RESEARCH
AND DEVELOPMENT



MARY JO FOLEY
ALL ABOUT
MICROSOFT
Editor

the JOURNEY CONTINUES

Join Microsoft & Industry Experts as they **Help you Navigate** the new and Exciting Technologies & Releases

EARLY BIRD DISCOUNT

Register by July 20 and save \$100 off the regular conference fee.



REGISTER TODAY! www.WinConnections.com • 800.438.6720 • 203.400.6121

Windows Server 2012

Is Good News for IT

With so many compelling features, there's a business case for most every company

I've spent the past few weeks presenting Windows Server 2012 community roadshows, a project sponsored by Microsoft where MVPs provide an overview to local IT professionals of the updated OS's many new and enhanced features. It's been an interesting time presenting the product for people who work with different parts of the technologies, in a wide variety of company sizes and industries. Without exception, everyone I've gotten feedback from has gone into the seminar skeptical and come away impressed.

A conversation I had with a consultant colleague of mine sums up the reasons people are excited about Server 2012. My friend hadn't been paying much attention to the server OS and asked me to tell him the top three features. I recognized his approach to considering a new OS because I'd used it in the past; historically with OS upgrades, there were usually just a few big features that made it to the top of one's "must have" list. In the case of Server 2012, three just wasn't enough. I had to limit myself to a top ten.

1. Hyper-V Shared Nothing Live Migration. Allows you to move a virtual machine (VM) from one Hyper-V host to another, without any common shared storage such as a SAN to hold the virtual hard disk, as long as there's an Ethernet connection between them. This is an industry first.



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.



Email



Twitter

2. Hyper-V network virtualization. Greases the skids for VM movement between hosts on different subnets or even service providers because you no longer have to re-address the VMs every time you move them.

3. Storage Spaces. Storage virtualization that allows you to use a wide variety of low-cost disks configured in flexible storage pools that all server components take advantage of.

4. Server 2012 Hyper-V. The free version of Microsoft's hypervisor, which has every capability of the full product except Windows Server licensing. It makes a perfect Linux or VDI host OS, and it far outstrips its free VMware competitor in both features and capacity.

5. Hyper-V Replica. Allows you to easily provide disaster recovery for a VM by loosely replicating it to another host in the next rack—or the next state—with no extra hardware.

6. Virtualization-safe and more easily deployed Active Directory. Unlike previous OS versions, you can't damage a virtualized Active Directory (AD) instance by performing improper operations on it. This new capability also allows you to safely clone new domain controllers (DCs) from existing ones.

7. IP address management (IPAM). A long-overdue IPv4 and IPv6 address space manager with more than 40 features that also manages all your DHCP servers from one console. It's an in-the-box alternative to address management products that cost five and six figures.

8. Server Message Block (SMB) changes. Huge advances in the SMB protocol that provide the ability to use Windows file servers in ways you would never consider before, such as high-performing, fault-tolerant remote storage for Hyper-V, VMware, or SQL Server.

9. High-availability DHCP. Finally! Available right out of the box.

10. Dynamic Access Control. Provides a new degree of control over file server data partly thanks to the new support of claims in AD.

A lot of IT pros who attended the workshops wanted to know which of these features would be in different versions of the OS.

Unfortunately, I don't have the answer. But I suspect by the time you read this we'll all know.

Despite significant new features in previous OS versions, IT pros remain uncomfortable changing their user interface habits. Regardless of [Jeffrey Snover's best efforts to convince IT pros of the necessity for PowerShell](#), I found that less than 10 percent of my audience is using it. I think that most IT pros won't start using this scripting language until they're backed into a corner and have no choice. But once they learn it, they'll be glad they did, especially with an OS like Server 2012 that allows you to do everything with PowerShell.

Even though these roadshows were at the tip of the Server 2012 local education wave, there's already an unprecedented amount of documentation available on Server 2012 features. [Microsoft's understanding and troubleshooting documents](#) for various roles and features are thorough and detailed (include "Windows Server 8" in your search terms), and many virtual labs are available. In addition, [TechEd 2012 deep-dive sessions into many Server 2012 features](#) are available on MSDN Channel 9. *Windows IT Pro* is also covering Server 2012 announcements and specific features on a regular basis. See the Learning Path for links to specific articles.

I predict that Server 2012 will be the most rapidly adopted OS in Microsoft's history. Even if you aren't looking at building your own private cloud, the OS has such a compelling set of features, for so many scenarios, that it presents a strong business case for a wide range of industries and company sizes. The Hyper-V enhancements, in particular, provide companies the means to safely virtualize their infrastructure with more flexibility at a far lower cost than ever before. ■

InstantDoc ID 143562



Learning Path

For more information about new features in Windows Server 2012:

"Microsoft Hyper-V Server 2012 Release Candidate Now Available"

"How Windows Server 2012 Improves Active Directory Disaster Recovery"

"Windows Server 8 Storage Spaces"

"Windows Server 8 Active Directory Moves Forward"

"Understanding Windows Server 8 Hyper-V Networking Changes"

"Server Management in Windows Server 8"

"Windows IT Pro Insider: Windows Server 8 and Windows 8"

"Exploring Windows Server 8: Dynamic Access Control"

Physical & Virtual Management of the Data Center

The story of datacenter management is one of rapidly increasing complexity—and the need to deal with what has become a much more flexible and much less static computing infrastructure. As consolidation reduced the number of physical servers in the datacenter, virtualization increased the number of servers that needed to be managed. And along with server virtualization came the need for improved management of networking and storage infrastructures in order for businesses

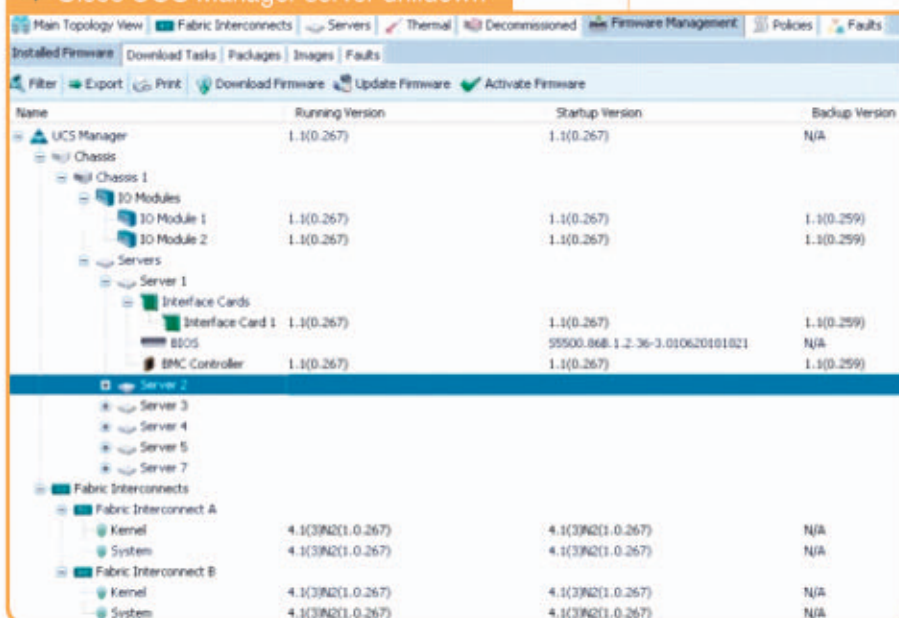
to derive full value from their now more flexible and adaptable data center.

As data centers move to a more unified infrastructure a key advantage is the ability for infrastructure vendors to deliver management tools that are completely integrated with all aspects of the data center. Management tools need to take into account the capabilities of the underlying infrastructure and allow data center administrators to monitor, manage, and automate control of data center behaviors

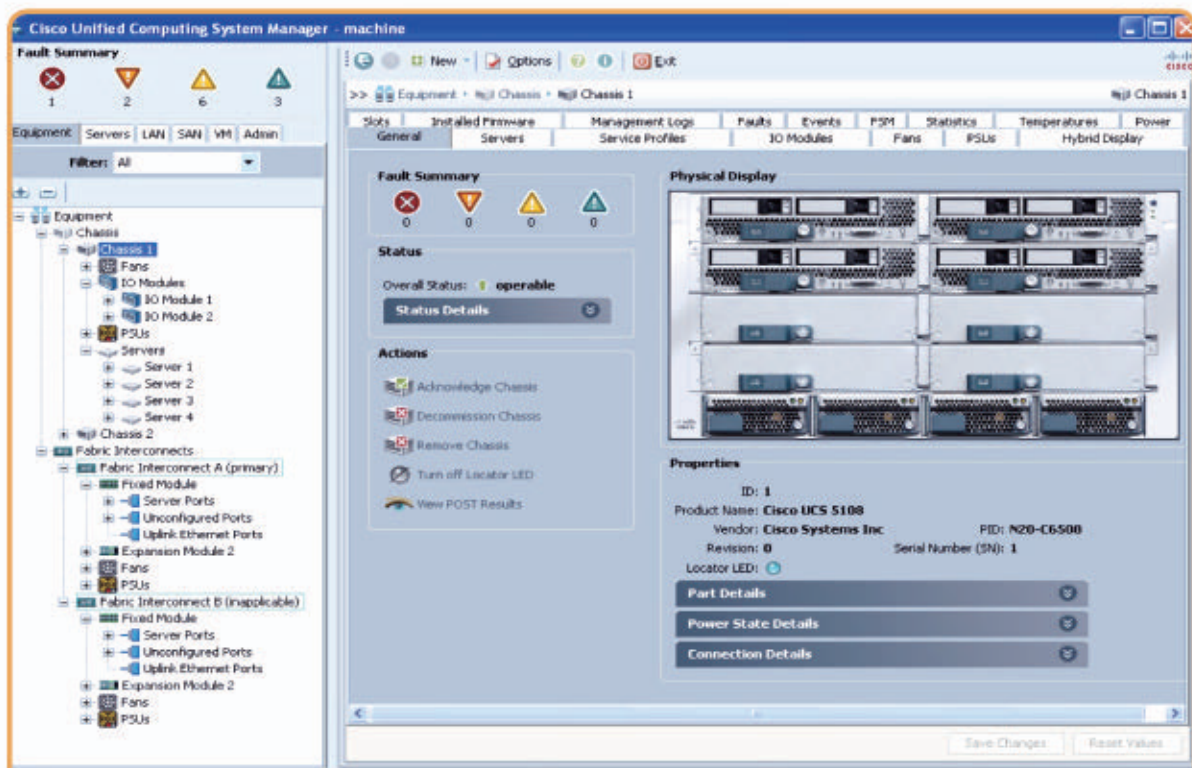
across servers, storage, and networking.

Tasks such as provisioning a new virtual server can require management actions across no less than three separate platforms. The virtual server must be created with the appropriate operating systems and server resources available, the network switch configuration may need to be

▼ Cisco UCS Manager server drilldown



Name	Running Version	Startup Version	Backup Version
UCS Manager	1.1(0.267)	1.1(0.267)	N/A
Chassis			
Chassis 1			
IO Modules			
IO Module 1	1.1(0.267)	1.1(0.267)	1.1(0.259)
IO Module 2	1.1(0.267)	1.1(0.267)	1.1(0.259)
Servers			
Server 1			
Interface Cards			
Interface Card 1	1.1(0.267)	1.1(0.267)	1.1(0.259)
OS		95500.868 1.2.36-3.010620101021	N/A
BMC Controller	1.1(0.267)	1.1(0.267)	1.1(0.259)
Server 2			
Server 3			
Server 4			
Server 5			
Server 7			
Fabric Interconnects			
Fabric Interconnect A			
Kernel	4.1(3P2)(1.0.267)	4.1(3P2)(1.0.267)	N/A
System	4.1(3P2)(1.0.267)	4.1(3P2)(1.0.267)	N/A
Fabric Interconnect B			
Kernel	4.1(3P2)(1.0.267)	4.1(3P2)(1.0.267)	N/A
System	4.1(3P2)(1.0.267)	4.1(3P2)(1.0.267)	N/A



▲ Cisco UCS Manager chassis view

changed, ports assigned or repurposed, routing and mapping may need to be addressed. And appropriate storage must be made available. You may need to make configuration changes across the network as access is given to those resources that will make use of the newly provisioned server, and the impact of all of the changes to the infrastructure needs to be assessed.

State-of-the-art infrastructure and datacenter management tools are able to integrate the management and operation of both physical and virtual infrastructure, giving IT the ability to see the data center environment as a comprehensive whole. But the needs of today's data center go far beyond simply being able to view

and report on the state of the infrastructure.

Complete data center infrastructure management requires the ability to do forward planning, what-if modeling, trending, and forecasting of future requirements—and that's just the start. Being able to implement changes in the computing infrastructure, automate the provisioning of new resources, and modify the networking infrastructure to properly support the addition, removal, or consolidation of resources are not just features that would be nice to have, but are absolute must-haves in a modern data center in order to get the greatest value from the corporate investment in the datacenter IT component. ●



BUILT FOR
THE HUMAN
NETWORK



YOUR BUSINESS NEEDS TO BE AGILE, FLEXIBLE AND RESILIENT. SO WHY IS YOUR SERVER ARCHITECTURE STATIC, COMPLEX AND OUTDATED?

When we designed our servers, we started fresh. No silos, no complexity. The result is a server unlike any other on the market. It's the Cisco Unified Computing System™. And it transforms efficiency and productivity. That's because Cisco UCS is based on simplicity, integration, speed, automation and ease. It's a difference our customers are noticing: 80% increase in administrator productivity. 90% reduction in deployment times. 40% improvement in application performance. 30% lower infrastructure costs. No wonder over 11,000 businesses have purchased Cisco UCS. It's built for productivity. Built for the future. Built by the only company in the world that could. Learn more at cisco.com/servers.

Cisco UCS is powered by the Intel® Xeon® processor.

Letters

letters@windowsitpro.com

Long Live Windows Phone!

Despite Paul Thurrott's gloomy prognostication—in his article “[Is Time Running Out for Windows Phone?](#)” (May 8, 2012)—I think a number of factors are working in favor of Windows Phone.

It's important to recognize that smartphones and computers don't share a marketplace and are driven by different purchase-and-replace decisions. Purchase cycles for smartphones, for example, are much shorter (typically 2 years) and they aren't burdened by entrenched infrastructure as are line-of-business (LOB) applications or even common applications such as Microsoft Word or Excel, which require extensive training and support. In fact, I'd argue that phones should be considered more like fashion purchases than technology and are more disposable. They're much easier to unplug and replace than, say, database applications. Overall, this market truth presents an opportunity for Microsoft, if only the company can take advantage of it.

Also, Windows Phone—in particular Windows Phone 8 (from what we know of it)—has distinct technical advantages (especially when compared with Apple iOS) that make it attractive to IT departments worldwide. For example, it can support multiple identities—an important trait as corporations increasingly allow and even encourage their employees to use personal devices for company work.

Another Windows ecosystem strength is its capacity to readily manage thousands of devices and the important identities inside them. If Microsoft extends this capability to Windows Phone 8, it should

Send Your Comments

Windows IT Pro welcomes feedback about the magazine. Send comments, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.



Comments



prove uniquely adept and highly capable of influencing corporate adoption. Even in an ordinary home environment, the inability of the iPad, for instance, to properly accommodate more than one user is a big failure. It shouldn't be necessary to have several tablets simply because different people use them. Microsoft can exploit this situation. Security also weighs heavily on the minds of IT (and increasingly consumers, too), and in this domain Windows also bests its competitors.

Of course, we can't understand the prospects for Windows Phone without also understanding the effect that Windows 8 tablets will have. These devices will likely offer a variety of functions and perhaps form factors that are currently unavailable through iOS (although I'm not certain whether this advantage will extend to Google Android) and which will have a positive overall effect on Windows Phone adoption. In particular, iOS's inability to properly handle styli is a big weakness: Note-takers, field workers, authors, artists, and others realize that Steve Jobs' stance in this regard was plain wrong. A pen or a stylus captures ideas in a vastly different way than fingertips and is quite a bit faster in many circumstances. I'm not saying that one is better than the other as much as I'm saying that one is preferred over the other for certain applications. Supporting both would be ideal, and it's an area of potential strength for Microsoft.

Many people point to the deliciously large number of iOS or Android apps available in their marketplace, and these can only be considered feathers in their respective caps. However, I personally use fewer than 10 apps, and I think this is true of most people—perhaps excluding gamers. Thus, after the main driver apps are re-created for Windows

or otherwise replaced by their competitors on the Windows Marketplace, this becomes a minor issue.

Although Mr. Thurrott makes valid points about the development environment, I think bemoaning the lack of an upgrade path for Windows Phone 7 devices is a bit unfair. People don't purchase clothes or furniture or stereo equipment expecting to be upgraded to the latest version when it comes, and that expectation doesn't have to be part of the smartphone paradigm either. As I said, they're switched out frequently enough to make it a non-issue. Yes, I'm disappointed that my Nokia Lumia 900 won't upgrade to Windows Phone 8, but if it hurts enough I can spend the \$250 to extend the contract and get a new device.

With all this noted—and generally being a Microsoft fan—I will concede that Windows Phone has a tough road to navigate. Microsoft has proven its mettle in the past and is capable of doing so again. ■

—Daniel Small

InstantDoc ID 143553

Windows and Visual Studio RCs and Office 15 Beta



Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for Windows IT Pro UPDATE, and a daily Windows news and information newsletter called WinInfo Daily UPDATE.

Email



Twitter



Website



While I focused exclusively on Windows 8 Release Preview last month, Microsoft was busy finalizing several key platform products, many of which are, or soon will be, available for evaluation. Everything is changing in 2012, and it all starts now.

Windows Server 2012 Release Candidate

Delivered alongside Windows 8 Release Preview in very late May, about a week ahead of the purported schedule, Windows Server 2012 Release Candidate (RC) is, if anything, even more mature and complete than its desktop sibling. Indeed, to the first point, Microsoft claims that Server 2012 (formerly code-named Windows Server 8) was, in fact, feature-complete as far back as the beta release in late February. And to the latter point, Microsoft revealed that its Bing search service was using Server 2012 RC for 100 percent of its worldwide queries, a stunning proof point for this OS's state of completion.

If you've been evaluating Server 2012, you know that this state of completion is, paradoxically, both factual and illusory. On the one hand, this Server release has always had an interesting vibe of solidity about it, a high-quality level that the Windows client team needed much, much more time to achieve. But on the other hand, Server 2012 is as much a slice in time as is Windows 8—something that, yes, Microsoft has to deliver on a predictable schedule but also needs to

update going forward since much of it seems incomplete, or at least less fully realized. And in the case of Server 2012, the incomplete bits, alas, include some key pieces of UI.

Chief among these is the new Server Manager, a strange Metro-like (but not “Metro-based”) dashboard that will cause some admins and IT pros fits until they get used to it. Server Manager succeeds in consolidating many of the tools and functionality that these folks will need to use every day, and it’s fairly successful as a front end to Server 2012’s new multi-machine management philosophy. But the weirdness of this UI, which is different from every other management interface in the system, is going to be off-putting to many.

Although admins and IT pros might be struggling with the interfaces in Server 2012, one thing they won’t quibble over is the technical capability of the system. We’ve come to expect fairly aggressive improvements in performance and scalability with each Windows Server release. But Server 2012 is particularly impressive in these areas.

For example, the maximum number of logical processors jumps from 64 in Windows Server 2008 R2 to 320 in Server 2012, an improvement of 500 percent. And maximum physical memory leaps from 1TB to 4TB.

On the virtualization side, the improvements are likewise incredible. The maximum number of virtual processors per host jumps from 512 in Server 2008 R2 to 2,048 in Server 2012. Virtual processors jump to 64, from 4. The maximum memory per virtual machine (VM) is now 1TB, up from 64GB. And the maximum number of active VMs leaps from 384 to 1,024.

These are the kinds of numbers that drive sales, since they allow for far higher workload densities than do previous Server versions. When combined with other new capabilities in Server 2012, particularly in Hyper-V 3.0, something tells me that this version of Server won’t be a tough sell at all, despite any qualms over the new UIs. You should evaluate Server 2012 RC for yourself, of course.

Visual Studio 2012 RC

Also delivered in tandem with Windows 8 Release Preview, Visual Studio 2012 Release Candidate (RC) provides a suitably near-final peek at Microsoft's next-generation developer tools and, as important, its developer platforms. Unlike Server 2012, however, Visual Studio 2012—formerly known as Visual Studio 11—has been rocked by some confusing last-minute changes. Actually, in that way, it's a lot like Windows 8.

You might not be surprised to discover that part of the problem is Metro. Visual Studio is of course a set of desktop applications, but in its mad bid to get everyone used to this style of user experience, Microsoft is also adding a “Metro-like” look and feel to many of its non-Metro applications, too. This includes Server 2012's Server Manager, as I mentioned, but also the applications in Office 15, and, as it turns out, the various versions of Visual Studio 2012.

When you hear the term “Metro,” you might immediately visualize flat tiles and tons of white space, but the Metro design philosophy is a bit more nuanced than that. For example, it pushes the notion of content, not application “chrome,” and the belief that the app's UI should get out of the way so that the point of the app—the content—can shine.

Unfortunately, this is where the Visual Studio team got things wrong, at least with Visual Studio 2012/Visual Studio 11 Beta: The resulting UI, designed to emphasize the code you're writing rather than the surrounding toolbars and widgets, was deemed an endless sea of green by developers, who complained en masse to the software giant about this dubious design choice. So one of the things Microsoft fixed in Visual Studio 2012 RC was to add color and contrast to the UI.

Developers weren't done complaining, however. In separate but related announcements timed closely to Visual Studio 2012 RC, Microsoft revealed which Visual Studio 2012 editions it would make available, including all of the expected paid versions (Professional, Test Professional, Premium, and Ultimate) as well as Visual Studio 11

Express for Windows 8 (with support for C#, C++, JavaScript, and Visual Basic—VB), Visual Studio Express for Windows Phone, Visual Studio Express for Web, and an as-yet-unnamed Visual Studio Express product for Windows Azure v.Next.

What was missing was a free Express product for the now-deprecated desktop, but Microsoft explained that amateur developers and students could continue to use the existing Visual C#, VB, and C++ Express 2010 editions instead. If you thought that was the end of it—after all, why would Microsoft support a legacy developer platform with brand-new tools when the existing ones work just fine?—you’d be very wrong. Free desktop development, like the color blue, is apparently a God-given right.

So I reacted with some amusement when Microsoft actually announced a few weeks later that it would create a new, free version of Visual Studio, called Visual Studio Express 2012 For Windows Desktop, that will support desktop-based C++, C#, or VB development. Why did Microsoft give in? I have no idea, but you can bet that after all the fuss, very few users will ever bother with this package anyway. I wish the Windows client team were this pliable. You can find the [RC versions of several Visual Studio 2012 editions](#) at the Microsoft website.

Windows 8 Release Preview Changes

Speaking of the client team, although I covered Windows 8 Release Preview last month, Microsoft has since revealed that it will be making several changes to the product after the Release Preview. I’ve argued that these changes suggest that the “feature-complete” tag on the Release Preview is a bit of a stretch, but Microsoft tells me that there are a few more important things to consider.

First, Microsoft routinely makes “fit and finish” changes to Windows as late as possible in the development cycle, which is to say between the final prerelease milestone (in this case the Release Preview, but normally the Release Candidate) and the final RTM version of the OS.

This is true enough—though with Windows 7, the changes were so minor that it’s perhaps forgivable that I’ve forgotten about this policy.

Second, much of what we consider to be Windows 8 is in fact decoupled from the normal, three-year Windows development time period. That is, the (Metro-style) apps that come with Windows 8 will be updated regularly moving forward and won’t need to wait for the next version of Windows—which is a major change from previous versions and a realization of the goal for the applications that used to be part of Windows Live Essentials. So even though Windows 8 Release Candidate ships with apps of varying quality, all of them will likely change before Windows 8 is finalized, then change again regularly over time.

These are both good points, but the sheer level of change we’re going to see post-Release Preview is somewhat alarming. The biggest change—the removal of the Aero desktop theme that Microsoft first previewed in late 2003, in favor of a flatter, more modern, and, yes, “Metro-like” theme—stinks of a last-minute flip-flop. Microsoft’s original goal for Windows 8 was that businesses would be able to roll out both Windows 7 and Windows 8, side by side, and that they wouldn’t need to retrain users because Windows 8 would look and work so much like Windows 7 that they’d be essentially identical.

This goal suggested that Microsoft would let businesses configure Windows 8 in such a way that users could boot directly to the desktop and skip the Metro-style Start screen. But the software giant has had a change of heart, and my sources at Microsoft tell me the company has been busy ripping out legacy code for the old Start button and Start menu so that developers won’t be able to write utilities that bring those features back. And boot to the desktop? Forget about it.

The move to a Metro-like desktop is obviously the final nail in the side-by-side usage coffin. Now, Windows 7 and Windows 8 will look absolutely nothing like each other, even when Windows 8 is used in desktop mode. And although it’s unlikely that Microsoft will ever admit this—the company is still pushing a pretty decent list of new

business-oriented features in Windows 8, too, of course—it's clear to me why this has happened: Businesses will never roll out Windows 8 anyway.

That's because they're already rolling out Windows 7 in record numbers. In fact, as Microsoft just revealed, it has sold over 600 million licenses for Windows 7 since October 2009. So Windows 7 is the next Windows XP, that version of Windows that will remain in use in businesses long after the software giant has moved forward to newer product versions.

Windows 8 will still go out on hundreds of millions of new consumer PCs and devices, sure. But most businesses will continue to downgrade, this time to Windows 7. So why bother with making Windows 8 look like its predecessor? (Microsoft's official story is that the Aero theme sucks battery life and that moving to a non-translucent theme will help. I guess I buy that, but it's not the full story.)

Coming Soon: Office 15 Beta and Windows Phone 8

There's more happening with Microsoft products, of course, and after Microsoft TechEd 2012 and a special Windows Phone Summit, I think it's fair to say we'll soon have a lot more to talk about with Office 15 (including new SharePoint and Exchange versions) and Windows Phone 8. To whet your appetite for this onslaught of newness, I'll tell you that I'm expecting to see the final release of Office 15 in early 2013, possibly accompanied—finally—by mobile app versions of key Office applications for [Apple iPad](#) and even [Android](#). Windows Phone 8, I'm told, will be based on Windows 8 and utilize a special version of the Windows RT developer libraries that's tailored for handsets. And the Visual Studio developer environment will run virtualized Windows Phone devices in . . . wait for it . . . Hyper-V 3.0. ■

InstantDoc ID 143397

Pipeline Problems and Get-ADUser

The pipeline is an essential tool for this cmdlet, but be sure you know how to wield it



**Mark
Minasi**

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.

Email



Twitter



Website



Last month, in “[More Control of Active Directory Groups through PowerShell](#),” I pondered how you might pick out just the users in the *folks* group who live in bigfirm.com’s Sales OU, and I imagined that you might have tried something like this:

```
get-adgroupmember folks | get-aduser -searchbase
    "ou=sales,dc=bigfirm,dc=com"
```

That would have been a very reasonable guess, given how you’ve used *get-aduser* to make up for *search-adaccount* and *get-adgroupmember*’s limitations. But it doesn’t work. Understanding why it doesn’t work will make you a more efficient PowerShell user, so let’s take that up this month.

First, let’s re-examine an example of the pipeline and *get-aduser* working well together from last month:

```
get-adgroupmember folks | get-aduser -pr lastlogondate
```

That showed you how to get past *get-adgroupmember*’s limitations vis-à-vis user properties. In the above example, the *get-adgroupmember* cmdlet retrieves all the accounts that are members of the *folks* group and sends them as Active Directory (AD) objects to the *get-aduser*

cmdlet via the pipeline, a concept symbolized by the “|” character. Pipelines let you combine two or more of PowerShell’s hundreds or thousands of narrow-focus cmdlets into a “one-liner,” a quickly constructed tool that gets useful work done. (If the concept of the pipeline is still unclear, imagine an assembly line with multiple stations. Each cmdlet is a station, like the part of an automobile assembly line where the brakes get installed or where the windshield is affixed.) Here’s a three-cmdlet one-liner I used a few months ago that generated a table of locked-out users sorted by how long it had been since they last logged on:

```
search-adaccount -u -l | sort -pr lastlogondate | ft
name,lastlogondate
```

Here, *search-adaccount* generated all the locked-out user accounts and passed them as an array of objects to the cmdlet *sort*, which sorted them according to their *lastlogondate*, resulting in that same array of objects arranged in a different order, and then passed that to the *format-table* (*ft*) command, which displayed them onscreen in a (somewhat) attractive fashion. The beauty of this is that *search-adaccount*’s author needn’t include a *-sort* option, and neither *search-adaccount*’s nor *sort*’s author need worry about offering formatted table output. In a nutshell, that’s the benefit of a pipeline.

That only worked, however, because PowerShell cmdlets are built to accept—to welcome!—objects from a pipeline as input and usually to generate potentially pipeline-able objects as output. Some cmdlets, however, are more easygoing about using pipeline input than others. *Sort* pretty much accepts any kind of object, as long it has at least one familiar-looking property such as *Name* or as long as you tell it what to sort on. *Format-table* will display anything displayable for objects you give it. But *get-aduser* is a bit pickier. Look at its Help to see what it expects as an input, and you’ll see that it insists that you either invoke it as *get-aduser -filter something* (as we’ve been using

it) or as *get-aduser -identity identifier*, which lets you find out about any one object by specifying its SID, samaccountname, objectGUID, or distinguished name (DN). (There's also a third option, *-ldapfilter*, but we'll ignore it for now.) *Get-aduser -identity*, then, retrieves only one user according to SID, logon name, GUID, or DN. The *-identity* option doesn't take wild cards! That being said, *-identity* makes perfect sense when you realize that it's there to make *get-aduser* a good consumer of information proffered by a pipeline. A closer look at *get-aduser*'s Help reveals that if you invoke *get-aduser* without *-identity*, *-filter*, or *-ldapfilter*, the cmdlet assumes that you want *-identity*. Thus, I've seen that this works:

```
get-adgroupmember folks | get-aduser
```

It works because when *get-aduser* sees no parameters, it assumes it's in *-identity* mode. In that case, it needs to see a samaccountname, an objectGUID, a SID, and/or a DN, or it's stumped. For example, the *dir* command in PowerShell returns a set of objects representing the files and folders in the current folder. So, if you typed

```
dir | get-aduser
```

you'd get a whole lot of red error messages because poor *get-aduser* would find itself awash in things that lack a samaccountname, an objectGUID, a SID, or a DN. OK, there's a bit more to it than that, but basically if you go with the pipeline-as-assembly-line analogy, *dir | get-aduser* essentially turns *get-aduser* into Lucille Ball in the chocolate factory. So, the moral of the story is that PowerShell cmdlets use the pipeline—but not all the same way! Check Help carefully before depending on them. Next month, I'll take another approach to our group/OU selection process. ■

InstantDoc ID 143394

Windows Server 2012 Storage Enhancements

Storage enhancements propel virtualization capabilities

Microsoft's Windows Server 2012 (formerly code-named Windows Server 8) is due out before the end of 2012, and without a doubt it's one of the most significant releases of the Windows Server OS to date. The enhancements in Hyper-V promise to put Microsoft's virtualization platform on the same footing as VMware, the new PowerShell integration will enable automated management, and the new Server Manager provides multiserver management. Even with all these new features, the storage enhancements in Server 2012 might very well be the most important changes. Let's explore the top 10 storage enhancements in Server 2012.

10 Revised Chkdsk—One of the more subtle, yet practical, storage enhancements in Server 2012 is the revised Chkdsk utility. With Windows Server 2008 R2 and earlier, if you needed to run Chkdsk, it required a dedicated system and it could take hours to complete on a large volume of stored data. The Server 2012 Chkdsk has been split into two parts. A background process scans the disk; the volume is still online during this process. The second part takes the volume offline while it performs the fixes on problems that were just identified during the scan.

9 VHDX disk format—An important storage enhancement for virtualization is the new VHDX virtual disk format. The new VHDX format significantly extends the size of Virtual Hard Disks (VHDs) by



Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).



Email

providing support for VHDs with up to 16TB of storage. Previous versions of the VHD format were limited to 2TB.

⑧ Offloaded Data Transfer (ODX)—Server 2012 has built-in SAN integration with the new ODX technology. ODX can provide significant performance improvements for copy and move operations performed on a SAN by enabling the work of copying or moving the data to be offloaded to the SAN, thus bypassing the need for the Server 2012 OS to handle the data movement.

⑦ SMB Multichannel—The new Server Message Block (SMB) 2.2 protocol is one of the biggest changes in Server 2012, and one of its coolest features is SMB Multichannel. SMB Multichannel allows multiple TCP connections to be established over multiple NICs for a single SMB session, enabling bandwidth aggregation of the multiple NICs and multiple CPUs involved. The result is greatly improved performance, giving SMB access comparable performance to directly accessed storage.

⑥ SMB Transparent Failover—The SMB Transparent Failover feature is another new SMB 2.2 feature. SMB Transparent Failover enhances clustering. If a hardware or software failure occurs on a cluster node, all SMB clients can transparently reconnect to another cluster node with no downtime.

⑤ SMB Scale Out—Another feature in the SMB 2.2 protocol, SMB Scale Out uses Cluster Shared Volumes (CSV) to store file shares that provide simultaneous access through all nodes in a clustered file server. SMB Scale Out provides better utilization of network bandwidth and load balancing of the file server clients.

④ Thin provisioning—Thin provisioning has finally made its way to the core Windows Server OS. Thin provisioning optimizes storage utilization by maximizing utilization of existing storage and reclaiming

any unused space (aka trimming). Server 2012 allows storage to be allocated on a just-in-time basis—it can identify thin provisioned LUNs, as well as provide notifications for exceeding threshold and physical resource constraints.

③ **Data deduplication**—Data deduplication works at the volume level and stores more data in less physical space. Microsoft claims it provides optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization data. Data deduplication uses sub-file, variable-size chunking and compression to segment files into small (32KB–128KB) variable-sized chunks. Then it identifies duplicate chunks, maintaining just a single copy of each chunk. Redundant copies of the chunk are replaced by a reference to the single copy.

② **Storage Spaces**—Server 2012's new Storage Spaces provides a storage solution to organizations and implementations where SANs and NAS devices aren't affordable. With Storage Spaces, the storage is virtualized. The physical disks providing the underlying storage are abstracted from the storage management and aggregated together to be used as a single pool. Data redundancy can be performed automatically.

① **Resilient File System (ReFS)**—NTFS has been with us since the initial release of Windows NT back in 1993. Server 2012 breaks new ground by adding ReFS. ReFS can verify and autocorrect data, and the file system never needs to go offline. ReFS isn't a replacement for NTFS. In the initial release, you can't boot from ReFS, and ReFS doesn't support conversion from NTFS. ■

InstantDoc ID 143157

The new SMB 2.2 protocol is one of the biggest changes in Windows Server 2012, and one of its coolest features is SMB Multichannel.

Intelligent Storage Means Better Performance

The storage market has long been defined by two primary objectives: capacity and performance. Throughout the history of computing the storage market has seen continuous growth in the capacity of individual drives, with regular upgrades in the I/O performance of storage devices via changes to the drive interfaces. In addition, improved controller and host bus adapter technology, along with data architectures spread across

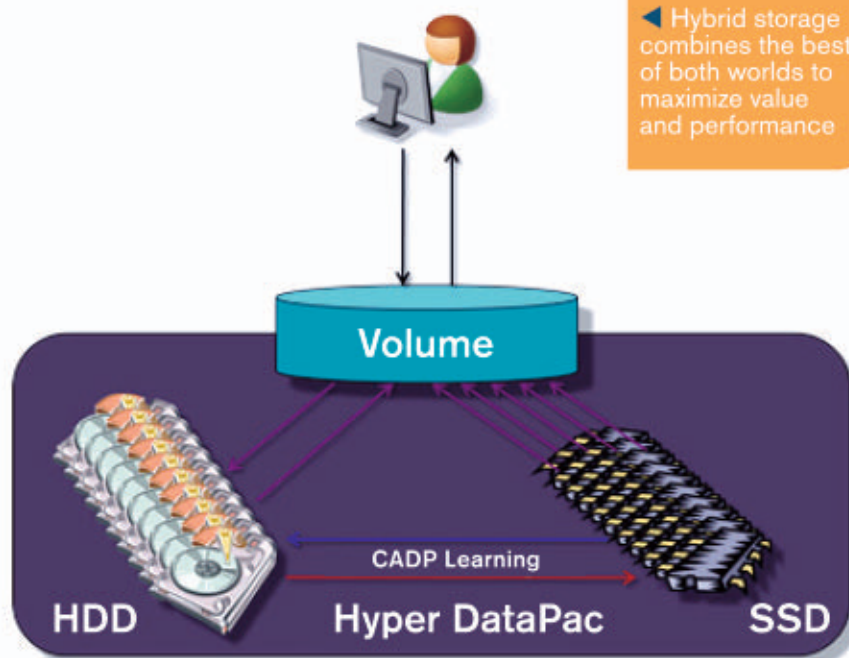
multiple drives, have continued to improve the I/O throughput of storage systems.


With the advent of Solid State Drive (SSD) storage a new technology has appeared in the form of hybrid devices and appliances. SSD has significant performance capabilities but at the cost of capacity and reliability. The most recent announcements by SSD manufacturers have focused more on improving reliability (e.g., Intel's announcement of five-year war-

ranties on certain SSD models) than on any revolutionary increases in capacity. The disadvantage that SSDs present in terms of capacity can be minimized or removed by combining expensive, low-capacity SSDs with inexpensive, high-capacity hard drives, which allows the faster SSD hardware to act as a front end for storage, in either SAN, NAS, or DAS format.

But this hybrid approach is not as simple as it might

◀ Hybrid storage combines the best of both worlds to maximize value and performance





▲ Intelligent Hybrid Storage eliminates performance drop off as capacity approaches full

appear. From the perspective of applications and servers that make use of the storage it is important that the hybrid approach look like any other high-performance storage device. Not only that, it is important to note that this isn't strictly a straightforward hardware solution; the storage needs to have a software/firmware layer that manages how data is migrated between rotating media and SSD storage.

Simply using SSD storage as a giant cache for the physical hard drives can provide some performance benefit, but this approach isn't cost effective, or the most efficient use of available storage. If that was the simplest way to do this effectively, an HBA with a giant cache would get you the same performance as a properly

designed storage appliance. The reality is that doesn't happen.

Technologies now exist that do a much better job of managing data on your storage systems. The performance and data availability issues start with applying well-established RAID technologies to the basic drive configurations, both for hard drives and SSDs, then extend simple striping and mirroring technologies to intelligent placement of data on the media best suited for it. Intelligent Hybrid Storage requires no manual configuration or policy to achieve high performance. User applications benefit from this performance throughout the lifecycle even while capacity approaches full, which is a key benefit from the investment of hardware. ●

It's everything

Your Business Needs
...And Wants.

- Unmatched Performance
- Ease of Management
- Xtreme Value



2011



2012

The award-winning performance of X-IO Hyper ISE has done it again. 2 years in a row, Hyper ISE has been recognized for ground-breaking performance and value like no other.

It's easy to see why X-IO continues to be the leader in data storage - **the fastest performance in the industry, exceptional reliability, up to 21.6 TB's of Storage in just 3U's of space, the lowest cost of ownership of any storage device** -- just to name a few.

The Hyper ISE combines the affordability of HDDs and the performance of SSDs with the best warranty in the business.

It's just smart business.



X-IO technologies

www.x-io.com

866.472.6764

Virtualization-Safe Active Directory in Windows Server 2012

AD gets a little smarter so that you don't shoot yourself
in the virtual foot

How many of you virtualize Active Directory (AD)? If you do, do you know and follow Microsoft's important guidelines about what you should and shouldn't do if you're virtualizing AD? That's what I thought. According to Microsoft's Customer Service and Support (CSS), AD is the top support area for Windows Server, and AD virtualization problems are at or near the top for AD issues. Shame on you! The good news is that Windows Server 2012's AD improvements will make your company's implementation safer for you slackers.

Unless you've been spending a lot of time spelunking in the [Cave of Kruber](#), as an IT person you're certainly aware of how virtualization has been taking over all aspects of the data center. The oldest—and thus most developed—area of virtualization is at the server level. Furthermore, because virtualization is one of the fundamental aspects of a computing cloud, its use is only accelerating.

If you're an AD administrator, what this means is that the virtualization team has been trying to get you to virtualize all of your AD domain controllers (DCs). AD administrators are by nature a risk-averse bunch (if we meet in person, I'll tell you the story of how I—OK, it wasn't me exactly—managed to instantly expire 30,000 Intel user accounts



**Sean
Deuby**

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.



Email



Twitter

during an upgrade, despite layers of precautions), but your first reaction was probably, “No!” However, these virtualization people are a tenacious bunch, and they probably want a good reason why they shouldn’t convert your entire AD forest to virtual machines (VMs). I defended my position (successfully) when I was at Intel. At the time, there were two good reasons, and although they still apply today, one of them is going away.

Reasons to Be Careful About Virtualizing AD

Security is the first reason not to virtualize. When I was pushing back at the virtualization folks, there were two aspects of security that we were concerned about: guest VM isolation and host administration. We were concerned about security between guests because we weren’t entirely convinced that one VM couldn’t gain access to another VM. As server virtualization has matured, that concern is pretty much put to rest. But the second reason—the concern surrounding operational security for host administration—is current and will be current for the foreseeable future.

Operational security for host administration is a long-winded way of saying that your virtualization host server administrators don’t necessarily know anything about the care and feeding of virtualized AD DCs. And in all current versions of Windows Server, server host or virtualization administrators can really screw up your AD installation if they use some of the basic capabilities that any virtualization product provides, such as image-based restores, rollbacks to snapshots, or virtual DC duplications.

The distributed nature of AD in Windows Server 2008 R2 and earlier can’t comprehend how virtualization products can change the state of a virtual DC in ways that can’t happen to a physical DC. And because it can’t comprehend and isn’t designed to handle these changes, the logical structure of this distributed system can lose its integrity, specifically in the form of USN rollback, an AD data integrity problem that’s difficult to detect and more difficult to recover from. Microsoft

has a comprehensive document about running virtualized DCs called “[Running Domain Controllers in Hyper-V](#),” and it includes a [section about USN rollback](#). If you suspect that you’re already a victim of a USN rollback, check the Microsoft article “[How to detect and recover from a USN rollback in Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2](#)” for information about how to detect and hopefully correct it.

How Server 2012 Makes AD Virtualization-Safe

Making AD completely safe in a virtualized environment was a top priority for the AD team, and they’ve achieved it in Server 2012. They haven’t just made it safe; they’ve enabled AD to take full advantage of virtualization’s capabilities. Conceptually, how it’s done is quite simple. First, you need to make a DC aware of when a rollback in time has happened. Second, the DC must take action that preserves its integrity and allows it to function normally.

To accomplish the first step, the layer that’s executing the change (the hypervisor) needs to flag that a rollback in time has occurred, and communicate it up through the virtualization stack. The application then has to recognize it. This process obviously requires that design changes be made to the hypervisor, the OS, and the AD application. The flagging mechanism is known as the VM-GenerationID.

The VM-GenerationID (or VM Gen ID) is a 128-bit value, held in the hypervisor, that represents the current generation of a VM’s state. As long as a VM continues to move forward in time, uninterrupted, the VM Gen ID doesn’t change. If the VM is rolled back in time—either from an image-based restore or from applying a snapshot—the ID is changed. This ID is mapped to an address in memory in the VM so that it’s available to applications running in the VM at all times. How does a DC know if its VM Gen ID has changed? When a Server 2012 DC is initially promoted (or upgraded), it stores the value of the VM Gen ID identifier in the msDS-GenerationID attribute on the DC’s computer object in its copy of AD during DC installation. Whenever

Technologically, Active Directory is fully virtualizable—but do you want to virtualize it entirely?

the DC reboots or processes a transaction (e.g., updating an attribute), it compares the current value of the VM Gen ID in memory with the value stored in AD. If they're different, the VM has rolled back in time and the DC must take certain measures to preserve integrity. The VM Gen ID is hypervisor-independent, and other hypervisor manufacturers (e.g., VMware) are building this capability into their products.

If a VM rollback has been detected, the DC performs two actions to prevent USN rollback: It resets the AD database's invocationID and dumps its local Relative Identifier (RID) pool. Resetting the invocationID (the version number of the local database) is the same action that's triggered if a normal restore process is run against the DC, and other mechanisms kick in to ensure that the DC has the latest updates from the other DCs it replicates with (including ones it created itself) but no longer has knowledge of due to the rollback. The RID pool is a collection of several hundred RIDs (part of the domain-unique SID) allocated to the DC by the RID master, to create SIDs when new security principals are created on the DC. Dumping the RID pool and requesting a new allocation from the RID master ensures that duplicate SIDs aren't created in the domain. Note that this doesn't get you out of regular backups, though!

So technologically, AD will be fully virtualizable, although as of this writing the Microsoft AD team hasn't quite decided whether they're going to make it official. But do you want to virtualize AD entirely? You must remember to look at the big picture before you decide. The modern data center has (or certainly will have, going forward) layers and layers of abstraction between the AD service and the all-too-fallible hardware. Remember the "all eggs in one basket" principle: Look at each layer below your service, work through the possible failure scenarios at each layer and how they'll affect your service, and plan your service configuration accordingly. For example, you should consider running more than one virtualization solution for some critical parts of your infrastructure so that a problem with one solution (e.g., a bad driver in the VMware ESXi kernel, or

an issue causing the Hyper-V parent partition to crash) isn't a single point of failure for your service. Similarly, even if your VMs are on different hosts and are using different virtualization products, if their virtual hard disks are stored on a single SAN it's a single point of failure. If the only way to mitigate one scenario is to have a few physical DCs, so be it! When the virtualization team objects, point out (probably to their second-level management) that the cost of maintaining a few physical boxes is trivial compared with the risk of your entire corporation being unable to log on in the morning.

Server 2012's Active Directory Domain Service (AD DS) has given you one less worry to keep you awake at night. But as with any new capability, you need to look at it in the context of your infrastructure and decide how you can best use it. ■

InstantDoc ID 143393

Troubleshooting Windows Server 2008 R2 Failover Clusters

Top things to look for



**John
Marlin**

is a senior support escalation engineer in Windows Commercial Technical Support, focusing on failover clustering. He has been with Microsoft for over 20 years. He is a Microsoft Certified Trainer for Clustering, delivering to internal Microsoft as well as Microsoft partners, and is a regular contributor to the [Ask the Core Team](#) blog.

Email



Blog



I want to discuss some of the troubleshooting techniques that we use with Windows Server 2008 R2 failover clusters. There are many ways to troubleshoot clusters, and some engineers might do things that others might not. So I wanted to pass along some of the most common things to look for and where to find them. With that in mind, let's first talk about the files that you'll generally be looking at and their descriptions.

One of the first things you'll be working with is Failover Cluster Manager, the new interface for managing a cluster. With this tool, you'll be managing groups and resources as well as performing some troubleshooting, which I'll explain as I go along. Failover Cluster Manager can be accessed from the Start menu and Administrative Tools.

Event Channels

First of all, you're probably familiar with the System event log. It's where we log critical, error, and warning events. However, it's not the only event log location that we write to. Starting in Server 2008, there are additional event *channels*. Figure 1 shows where to find the channels relevant to failover clustering. Here is where we'll log all the informational-type events and debug/diagnostic events. You'll find the following list of logs and their channels:

- FailoverClustering
 - Operational
 - Diagnostic (if *Show Analytic and Debug Logs* is selected)
 - Performance-CSV (if *Show Analytic and Debug Logs* is selected)
- FailoverClustering-Client
 - Diagnostic (if *Show Analytic and Debug Logs* is selected)
- FailoverClustering-Manager
 - Admin
 - Diagnostic (if *Show Analytic and Debug Logs* is selected)
- FailoverClustering-WMIProvider
 - Admin
 - Diagnostic (if *Show Analytic and Debug Logs* is selected)

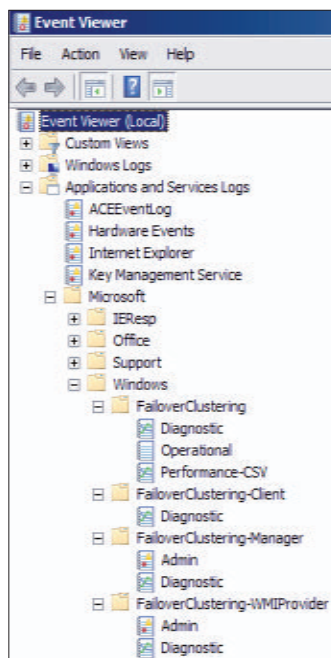


Figure 1
Channels relevant to failover clustering

If you're starting/stopping the cluster service, or you're moving groups, or groups are coming online and offline, and so on, those events will be logged in the FailoverClustering\Operational log. For example:

Event ID: 1061

Description: The Cluster Service successfully formed the failover Cluster "JohnsCluster"

Any failures connecting to other nodes opening Failover Cluster Manager are logged in FailoverClustering-Manager\Admin. For example:

Event ID: 4684

Description: Failover Cluster Manager could not contact the DNS Servers to resolve name "W2K8-R2-NODE2.contoso.com". For more information see the Failover Cluster Manager Diagnostics channel.

If you look at the FailoverClustering-Manager\Diagnostic log, you would see this:

Event ID: 4609

Description: An error was encountered while attempting to ping "W2K8-R2-NODE2.contoso.com". System.ApplicationException: Could not contact one or more DNS Servers. Please verify that DNS configuration is correct and the machine is fully connected to the network.

Event ID: 4612

Description: Server W2K8-R2-NODE2.contoso.com ping failed.


Just from these events, you can see that there is a problem with the node getting to the DNS server and can start troubleshooting this specific problem. What you might see without looking at these logs is possibly the W2K8-R2-NODE2 showing as down in Failover Cluster Manager. (One of the other logs mentioned above is the Failover-Clustering\Diagnostic log. I'll discuss this log a bit later.)

Failover Cluster Manager

To make things a bit easier, you can also view system event errors and warnings from within Failover Cluster Manager. On the main page in the middle pane, there is a Recent Cluster Events link that you can select, as Figure 2 shows. This link provides a handy way to display all warnings and errors that have occurred with Failover Cluster as the source in the past 24 hours. It pulls these events from all nodes and gives you everything in one spot. So there's no need to go to multiple machines and have multiple event logs open that you must switch between.

Figure 2

Recent Cluster Events



You can use the Query option to look for specific events. On the main page in the left pane, you'll see Cluster Events. You can right-click

Cluster Events and choose Query, or you can select Query from the Actions pane on the right. Figure 3 shows the Cluster Events Filter.

This is also a good way to display everything in the same location. For example, suppose you're experiencing the failure of a disk resource. You can bring up Failover Cluster Manager and have it query all nodes, the System event log, the error, and the specific date. On the main page, you can see when the disk failed, on what node(s) it failed, and any other pertinent data (such as disk events where a path failed).

You also have the ability to save these queries for later use.

You have two more options for looking up events. You can look up all resource-failure events for anything in a group, or you can be resource-specific. In the Actions menu, which Figure 4 shows, you can select *Show the critical events for this application* (any resource in the group) or *Show the critical events for this resource* (only the specific resource). Doing so will bring up the query for any of the events in the current event logs on all nodes. This option can also be beneficial for determining history and whether the event can be narrowed down to a specific time period or node.

For those who remember the Windows 2003 Server Cluster days, this is the Cluster.Log equivalent. Starting in Server 2008 failover clustering, the functionality is more in line with the Event Tracing for Windows (ETW) process. Instead of writing to a Cluster.Log text file, it writes it to a Diagnostics log located in the C:\Windows\System32\winevt\logs folder. There are three diagnostics logs that

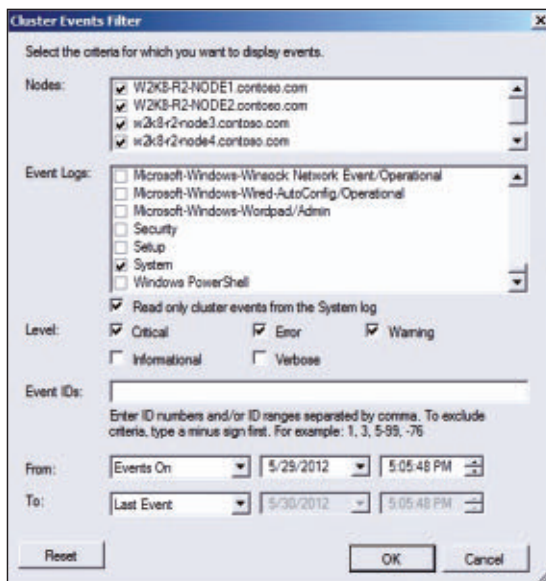
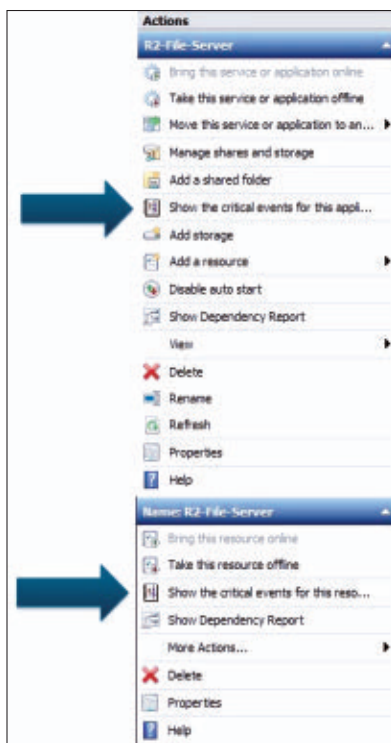


Figure 3
The Cluster Events Filter

Figure 4
The Failover Cluster
Manager's Actions
menu



we write to (clusterlog.etl.001, clusterlog.etl.002, and clusterlog.etl.003). We're only going to write to one of these at a time on any given boot. For more information about these log files and how they're used, check out the "[Understanding the Cluster Debug Log in 2008](#)" blog post.

This log is enabled and always writing. If you right-click FailoverClustering\Diagnostic and select *Disable log*, you can see all the events it has written. If you disable this log, the system will no longer write to it and information won't be saved. If you do this, it's best to save the event out as an event log or text file and enable it again. There are essentially three main events you'll see:

- Event 2049 is an informational event.
- Event 2050 is a warning.
- Event 2051 is an error.

These events will only be from the current diagnostic Event Trace Log (ETL) being written to. You'll see the event information just as you would the System or Application event log. However, each event will be only one line at a time. So, going event by event through this diagnostic event log can be pretty tedious. You can create a Cluster.Log text file with commands that combine all three of these logs into one to make the review of it much easier.

The PowerShell Get-ClusterLog command goes out to all nodes and generates a Cluster.Log on each node and places it in the C:\Windows\Cluster\Reports folder. This would be the Cluster.Log you might be more familiar with from Windows 2003. There are Get-ClusterLog

switches you might want to consider, depending on the circumstances. For example, say you can reproduce a failure at will and need to find the reason for the failure. Simply reproduce the problem and use the command

```
Get-ClusterLog -TimeSpan 5
```

to get data from the past 5 minutes. Because you need only the log from the one node you reproduce the problem on, you could add the *Node Nodename* switch to create the Cluster.Log on this single node. If you have a number of nodes and need to send these logs, it might take some time to connect to each node to get the file. In these circumstances, you could use the *-Destination* switch. This switch creates a Cluster.Log for each node, copies it to a folder you specify, and tags the name of the machine as part of the file name (e.g., W2K8-R2-Node1_Cluster.Log).

Remember that the Cluster.Log you're creating is a snapshot in time. It will take what's there right now and won't update with anything after it's generated. When it's generated, if there's a Cluster.Log in the Reports folder, it will get deleted to make room for the new one.

Resource Host System

The next thing I wanted to discuss is the Resource Host System (RHS). One of its responsibilities is to monitor the health of all resources in the cluster. It does this through a series of checks (basic and thorough). If a resource doesn't respond to these checks, RHS will issue the following system event:

Event ID: 1230

Description: Cluster resource 'Cluster Disk 1' (resource type '', DLL 'clusres.dll') either crashed or deadlocked.

In this instance, the disk didn't respond that the health check was made. What the cluster will do is fail the resource and restart it to get

you back to production. If these checks weren't in place, it could lead to a hung machine or no connectivity from a client application.

When troubleshooting the RHS event, you must consider the resource. If a disk deadlocks, you would need to consider everything in the disk stack. Was there slow disk I/O? Did you lose a path to the drive? This would be the focus of your troubleshooting. So, next up is reviewing the System event log for disk-related events, looking at Performance Monitor, updating drivers, and so on. If the resource was an IP address or a network name, your focus would be the network stack and everything there.

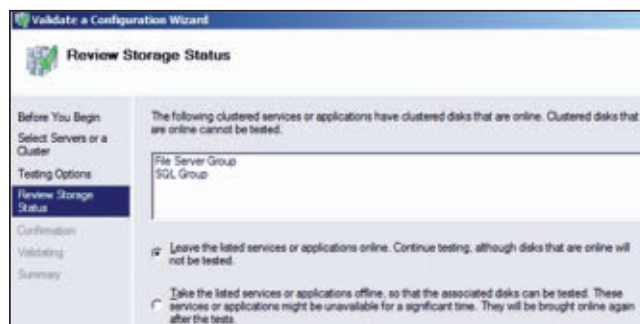
Cluster Validate

The last thing I want to mention is the Cluster Validate report. For a cluster to be “certified,” all components must be listed on the Windows Server Catalog, and it must pass a full Cluster Validate. Many people will run Cluster Validate before the cluster is created or just after. However, if there is a problem later on, few people remember to run Cluster Validate. You can use it as a troubleshooting tool! If you're having some disk problems, run the Storage Tests. If you're having network-communication problems, run the Network Tests. You can also use Cluster Validate to get information about groups, resources, and settings for your currently running failover cluster to be referenced at a later time.

The nice thing about Cluster Validate is that you can run it even while in production. When you run it and select the Storage Test, it will ask if you want to take the running groups offline, as you

see in Figure 5. The default setting is to leave the online groups alone, so production won't be affected. For the Storage Tests, it will test disks that are:

Figure 5
Running Cluster
Validate



- In groups that are offline
- In the available storage group
- Not a part of the cluster

Each time you run Cluster Validate, it will create a file in the C:\Windows\Cluster\Reports directory and will tag the date and time as part of the file name. So, every time you run it, it will create a new one and will create the file on all nodes that Cluster Validate was run against.

There are other ways to troubleshoot failover clusters—I just don't have enough space to cover them all. However, this column should get you started for most of the problems you might face. For more information, check out the [Ask the Core Team](#) blog and the [Clustering and High Availability](#) blog. Happy clustering! ■

InstantDoc ID 143535



The advertisement features a dark blue background with a faint world map. In the top left corner, there is a green and blue globe showing the Americas. The main text is in large, bold, white capital letters. The title 'CREATING INTERNET FACING SHAREPOINT SITES WITH SHAREPOINT 2012' is positioned in the upper half. Below it, 'SEPTEMBER 6TH' is written in a very large font. Underneath that, 'FREE VIRTUAL CONFERENCE' is followed by a vertical line and 'SPEAKERS INCLUDE TODD BAGINSKI, CELINA BAGINSKI AND MORE!'. At the bottom, a white rounded rectangle contains the text 'CLICK ON THIS AD TO REGISTER AND FOR MORE INFORMATION!'.

**CREATING INTERNET FACING
SHAREPOINT SITES WITH
SHAREPOINT 2012**

SEPTEMBER 6TH

FREE VIRTUAL CONFERENCE | SPEAKERS INCLUDE TODD BAGINSKI,
CELINA BAGINSKI AND MORE!

**CLICK ON THIS AD TO REGISTER
AND FOR MORE INFORMATION!**



Mike Danseglio



Jan De Clercq



William Lefkovich



Avril Salter



John Savill



Greg Shields

FAQ

Answers to Your Questions

Q: Can I create a software RAID volume by using Windows Server 2008 R2 on a USB-connected drive?

A: To create a software RAID volume such as RAID 1 (mirroring) or RAID 5 (striping with parity), the disks must be converted from basic to dynamic. Dynamic disks aren't supported on any kind of detachable storage, which includes USB drives, as described in the Microsoft TechNet article "[Change a Basic Disk into a Dynamic Disk](#)." This means software RAID volumes can't be created by using USB-connected media in Server 2008 R2 or any earlier OSs.

—John Savill

InstantDoc ID 143164

Q: Can I create a software RAID volume on a USB-connected drive in Windows Server 2012 Storage Spaces?

A: Server 2012 (formerly code-named Windows Server 8) provides a new capability called Storage Spaces, which allows inexpensive storage to be pooled together to create a storage pool. Volumes, or spaces, such as a striped disk, a mirrored disk, or a stripe with parity can then be requested from that pool. As a user, you don't need to

know which disks are being used from the pool—it's all performed automatically. If a disk fails, Storage Spaces automatically uses another disk in the pool, if available, to repair any volumes affected. The good news with Storage Spaces is that the types of storage have been increased from Windows Server 2008 R2 dynamic disks and now include USB-connected storage, in addition to SATA and Serial Attached SCSI. A storage pool can contain a mix of connectivity, mixed types of storage, and even different sizes. So the short answer to the question is, Yes!

—John Savill

InstantDoc ID 143164

Q: Can I use Managed Service Accounts (MSAs) on my Windows 2008 R2 servers if my domain isn't at Windows Server 2008 R2 domain mode?

A: Managed Service Accounts (MSAs) allow accounts to be created in Active Directory (AD) whose passwords are automatically synchronized and updated per policies with the server that uses the accounts for a service. This avoids the common problem of using accounts that either have to be set with passwords that never expire, which is a security risk, or having passwords that expire and break the service.

Although the domain functional level doesn't need to be Server 2008 R2, the Server 2008 R2 schema update must have been applied to the forest. If the domain isn't at Server 2008 R2 functional level, the Service Principle Names will need to be manually managed (but passwords are still handled automatically).

—John Savill

InstantDoc ID 143274

Q: How can the security of a Windows service benefit from the service isolation feature? How can I set up service isolation for a given Windows service?

A■ Service isolation lets administrators control which local resources a Windows service can access (e.g., files, registry keys on the local machine). Microsoft introduced service isolation in Windows Vista and Windows Server 2008.

On previous Windows versions, when you use one of the built-in high-privilege local accounts for running the service (Local System, Network Service, or Local Service—this is the account you configure on the Log On tab of the service’s properties) and you grant that service account access to a resource (in the access control settings of the resource), you also implicitly allow all other services that run under the same service account to access that resource, whether they need access to it or not.

To avoid this security problem, administrators typically create a custom Windows account to run their services. This solution, however, creates additional account management overhead. When you create custom service accounts, you also can’t leverage the automated password management features that Windows provides for the built-in accounts.

Service isolation builds on new service-specific SIDs that you can enable for each Windows service. These SIDs let administrators isolate a resource for a service’s exclusive use by securing the resource with an access control entry (ACE) that refers to the service’s service-specific SID.

A service-specific SID is linked to the service’s name (e.g., My-service), and not to the service account (e.g., LocalSystem). Service-specific SIDs are used for authorizing a service; the service account is used for authenticating a service.

Thanks to service-specific SIDs, administrators don’t need to worry anymore about the cumbersome process of creating and maintaining custom service accounts. They can now continue to use the built-in service accounts for authenticating their services while relying on service-specific SIDs for authorizing services and setting permissions on the services’ resources.

To generate a service-specific SID for a service, you must set the service's SID type property to either "unrestricted" or "restricted." You can do so by using the `sc` command:

```
sc sidtype unrestricted
```

After you run this command, the next time the service starts, Windows adds the service-specific SID to the access token of the process that hosts the service. Service-specific SIDs are local (machine-level as opposed to domain-level) SIDs.

To query the current value of a service's SID type property, you can use the following `sc` command:

```
sc qsidtype
```

After you've enabled a service-specific SID for a service by using the `sc` command, you can leverage the SID to restrict access to the service's local file system and registry resources. To refer to a service-specific SID when setting permissions on resources, you must use the syntax `NT SERVICE\`.

Windows constructs service-specific SIDs from the service name by using the following formula:

$$S-1-5-80-\{\text{SHA-1}(\text{service name in uppercase})\}$$

where SHA-1 stands for the application of the SHA-1 hash function on the service name in uppercase. This formula illustrates how, for a given service, the service-specific SIDs will be identical on different Windows machines. This system greatly facilitates resource permissioning for an identical service that might be hosted on different Windows machines.

—Jan De Clercq
InstantDoc ID 143215

Q: What is context-aware computing, and how might it aid healthcare IT?

A: Many of the technologies we take for granted today have their origins in solving the unique problems found in healthcare IT. If you think about the paths doctors take as they're running between examination rooms on a very limited time schedule, you'll quickly see how their every technology interaction must be highly optimized.

Presentation and desktop virtualization technologies have greatly assisted this process by enabling doctors to disconnect and reconnect to remote sessions without waiting for logons and logoffs. Yet even a highly optimized [XenApp](#) or [XenDesktop](#) infrastructure can't dynamically reconfigure a doctor's user session at each connection without extra help. Recognizing each new situation and adapting that session's configuration (i.e., running applications, connected printers and devices, and so on) are what context-aware computing solutions add to XenApp and XenDesktop.

Many vendors for these solutions tend to specialize in a specific vertical. The Aventura HQ solution from [Aventura](#), for example, is specifically designed for use in healthcare situations. Other solutions are more generally applicable, such as [RES Software's](#) Dynamic Desktop Studio.

Not every environment requires the level of situational relevance that these products enable; however, if your users need to connect from multiple locations throughout the day with location-relevant settings, context-aware computing solutions such as these can greatly improve that user-to-computer interaction.

—Greg Shields

InstantDoc ID 142527

Q: What is 4G? Does it mean different things to different industries?

A■ There are two definitions for 4G—a technical definition and a marketing definition. The technical definition for 4G is that it's a technology that meets the International Telecommunications Union's (ITU's) IMT-Advanced specifications. These specifications include a data rate of 100Mbps at low mobility pedestrian speed. The technology LTE-Advanced meets these ITU requirements and is therefore 4G.

Many service providers are deploying technologies that offer data rates higher than the initial 3G systems but not high enough to meet the ITU IMT-Advanced requirements. These technologies include WiMAX, HSPA+, and LTE. To distinguish these technologies and help sell the high-speed capabilities, service providers are marketing these services as 4G.

—Avril Salter, Mike Danseglio

InstantDoc ID 142417

Q■ What can I use keyboard shortcuts for in Microsoft Outlook?

A■ Microsoft provides its Windows OSs with many keyboard shortcuts to increase user efficiency and provide alternative input mechanisms for basic program functionality. Like most major applications, Microsoft Outlook uses Windows keyboard shortcuts for specific Outlook functionality and interface manipulation.

Microsoft offers hundreds of keyboard shortcuts for Outlook. Many of the shortcuts are navigational, such as moving between or jumping to a specific folder (e.g., Ctrl + 1 for Mail, Ctrl + 2 for Calendar).

One keyboard shortcut that has always been mildly confusing to me is the one to initiate a *Send and Receive* in Outlook. I consider the *Send and Receive* command a form of refresh for email folders. In many Microsoft applications, F5 is commonly used to refresh the current interface and its rendered content. But in Outlook, F9 has always been the shortcut key to launch a manual *Send and Receive* request.

Here are a few keyboard shortcuts I use in Outlook. These shortcuts aren't as common as, say, Alt + S to send a message or Ctrl + N to create a new item.

Wherever you are in Outlook, you can easily launch the New Search Folder window by selecting Ctrl + Shift + P. I prefer keyboard shortcuts I can select with one hand, such as Ctrl + C to copy and Ctrl + V to paste. I use very few that need three keys and both hands; however, I do search folders a lot and this keyboard shortcut is one I favor.

When you're within the Calendar view in Outlook, the interface offers Day, Work Week (in Outlook 2010), Week, and Month views of the calendar. You can also get a 10-day view with a simple keyboard shortcut. With a day selected in the Outlook calendar, it's a simple Alt + 0 to render a 10-day view starting with the day that was highlighted. In fact, you can view any number of days from 1 through 10 by using Alt plus the number of days you want to see. For example, Alt + 8 shows a view of 8 days in your Outlook Calendar. If you're displaying multiple calendars, this setting applies to all of them simultaneously.

If you code simple macros or work in Visual Basic for Applications (VBA) through Outlook, you probably use Alt + F11 to launch the Visual Basic Editor within Outlook. The Visual Basic Editor opens the same file (if any, and if available) that was accessed when it was last closed.

Microsoft publishes a formal list of keyboard shortcuts for different versions of Microsoft Outlook: see [“Keyboard shortcuts for Outlook 2010,”](#) [“Keyboard shortcuts for Outlook 2007,”](#) and [“Keyboard shortcuts for Outlook 2003.”](#) You can expand the categories of shortcuts and print their pages for a complete reference.

—William Lefkovich
InstantDoc ID 141448

Q: How do I use Windows PowerShell to switch between Windows Server 2012 configuration levels?

A■ Using the `Uninstall-WindowsFeature` and `Install-WindowsFeature` cmdlets makes it easy to add and remove the components for the configuration level you want to use in Server 2012. For a more extensive explanation about switching configuration levels, see my FAQ [“Is it true Windows Server 2012 allows a server to switch between Server Core and Full Install mode without reinstalling the OS?”](#)

The two features you’ll be adding or removing are `Server-Gui-Mgmt-Infra` for the management tools and `Server-Gui-Shell` for the graphical shell. It’s not required to specify the `Server-Gui-Shell` to be removed; removing the `Server-Gui-Mgmt-Infra` forces the removal of `Server-Gui-Shell` because it’s dependent. To take a full server and make a Server Core installation, use the following command:

```
Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -Restart
```

To take a Server Core installation and make it a full server with a GUI, use this command:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra,  
Server-Gui-Shell -Restart
```

Likewise, you can add just `Server-Gui-Mgmt-Infra` to make a minimal server installation. Or you can add `Server-Gui-Mgmt-Infra` to a Server Core installation to make a minimal server installation:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra -Restart
```

—John Savill

InstantDoc ID 143249

Navigate VMware Licensing

Don't let licensing changes blind you to vSphere 5.0's top 5 features



Alan Sugano

is the president of ADS Consulting Group, which specializes in virtualization, networking, custom programming, Microsoft .NET web development, and SQL Server development. He's the author of *The Real-World Network Troubleshooting Manual* (Charles River Media).



When VMware released vSphere 5.0 in August 2011, there was more discussion of the licensing changes—specifically, virtual RAM (vRAM) entitlement—than of the new features. But in my environment we've migrated about half our clients to vSphere 5.0, and it appears stable. The initial issues with backup support have been resolved; there are workarounds to get a proper virtual machine (VM) backup image running on vSphere 5.0. What was all the licensing fuss about, how can you best avoid it, and which features make it all worthwhile? Read on to find out.

vRAM Entitlement

One of the more controversial issues with vSphere 5.0 is the addition of a memory entitlement model. When vSphere 5.0 was first released, the memory entitlements were significantly more restricted than in earlier versions. Since then, because of feedback from the VMware Community, VMware decided to increase the amount of memory entitlement that comes with each socketed CPU.

With vSphere 4.x, licenses had some limitations, based on the number of cores in the CPU. Of course, with the current generation of servers, the number of cores per CPU has increased significantly. You can now purchase a CPU that has 10 cores! With vSphere 5.0, VMware made a move to license VMware ESX based on the amount of vRAM that you get with each CPU socket license. Table 1 shows the amount of vRAM that you get with each version of vSphere.

For example, if you have a two-socket VMware ESX host running vSphere Enterprise, then you have 128GB (i.e., two sockets at 64GB each) of vRAM entitlement on that ESX host.

Table 1: vSphere Editions and vRAM Entitlements

vSphere Edition	vRAM Entitlement
Essentials	32GB
Essentials Plus	32GB
Standard	32GB
Enterprise	64GB
Enterprise Plus	96GB

This means that you can run VMs that have a total requirement of as much as 128GB of memory on the ESX host. Your 12-month rolling average memory utilization must be equal to or lower than your vRAM entitlement. If you exceed this limit, VMware won't prevent a VM from running, but you'll receive VMware vCenter Server alerts informing you that you're out of compliance.

You can also pool your vRAM entitlements. For example, suppose that you have a two-host ESX cluster running vSphere Enterprise Plus. In this configuration, you need a minimum of four sockets of Enterprise Plus, giving you a total of 384GB (four sockets at 96GB each) of vRAM entitlement. It doesn't matter whether all the VMs are running on one host and the other host is idle, as long as the total configured memory on all the VMs does not exceed 384GB.

The good news (sort of) is that the maximum amount of vRAM is capped per VM at 96GB. Therefore, the most that a single VM can count against the vRAM entitlement pool is 96GB, regardless of how much memory is configured for the VM. So a VM that's configured

with 512GB of memory consumes only 96GB of the vRAM entitlement pool. Put another way, the vRAM entitlement requirement for three 1TB VMs is 288GB (three at 96GB each).

What if I Need More vRAM?

If you start to receive warnings that you have exceeded your vRAM entitlement, you have several options for increasing the vRAM entitlement pool:

- Purchase additional licenses of the same vSphere edition.
- Upgrade the existing vSphere edition to a higher level. Going from vSphere Standard to vSphere Enterprise gives you an additional 32GB of vRAM entitlement per socketed CPU license. Of course, if you have Enterprise Plus, this option won't work: Enterprise Plus has the largest vRAM entitlement of all vSphere editions.
- Introduce another host (or hosts) running the same vSphere edition as the existing hosts in an ESX cluster. Be aware that if you want to maintain VMware vMotion compatibility, you'll need to get a host that has identical or nearly identical CPUs.

Note that you can't extend the amount of vRAM in the vRAM pool for vSphere 5.0 Essentials or Essentials Plus. Only vSphere 5.0 Standard through Enterprise Plus editions can be extended indefinitely, and you must purchase the same edition of vSphere to expand an existing vRAM entitlement pool.

If you used vSphere 4.x to deploy a virtual desktop infrastructure (VDI) solution, purchased your vSphere licenses before September 30, 2011, and have a valid Support Agreement, then you can upgrade to vSphere 5.0 while retaining an unlimited vRAM entitlement pool. However, you must use a separate instance of vCenter Server that manages only the VDI. Any vSphere licenses that you purchase separately to run VMware View (VMware's VDI solution) are still subject to the vRAM Entitlement licensing model. (See the VMware Community post [“Desktop Virtualization with vSphere 5: Licensing Overview”](#))

for more information.) And there's good news for vSphere Advanced edition owners: If you were fortunate enough to purchase vSphere 4.x Advanced and have a valid support agreement, then you're entitled to vSphere 5.0 Enterprise. There's no Advanced version of vSphere 5.0, so you are grandfathered in to vSphere 5.0 Enterprise.

Top 5 vSphere Features

Significant changes in vSphere 5.0 can improve manageability and increase the return on investment (ROI) of your virtualization infrastructure investment. After you've got the licensing cleared up, take a look at these top new features of vSphere 5.0 and how we use these features in different environments.

#1: ESXi only. A few years ago, VMware announced that vSphere 4.1 was the last version to include both ESX and VMware ESXi. Keeping that promise, vSphere 5.0 comes only with ESXi, which basically is ESX without the Red Hat-based Service Console and the Web Server. ESXi has a significantly smaller footprint and is more secure than ESX. However, if you have any existing programs (e.g., a backup agent) that run in the Service Console, then you'll need to find a replacement for these services when you upgrade to vSphere 5.0.

For many of our clients, we were running Symantec Backup Exec's Remote Agent for Linux and UNIX Servers (RALUS) to obtain image backups of the VMs that were running on the host. To move to vSphere 5.0 and still get image backups of the VMs, we needed to purchase a vSphere 5.0 backup solution, such as [Veeam Backup & Replication](#), [Quest Software's vRanger Pro](#), or [Symantec Backup Exec](#). Note that these backup solutions require you to back up to disk, so you can perform a granular restore of a single file by using the VM image backup file. This requirement might mean that you need to purchase additional hard disks or a NAS appliance so that you have enough room to back up the VMs to disk. Of course, we still recommend moving the image backups to some type of offline media (e.g., tape) in case you are hit with a nasty virus that could wipe out your VM-to-disk backup files.

Significant changes in vSphere 5.0 can improve manageability and increase the ROI of your virtualization infrastructure investment.

#2: Support for as much as 1TB of memory on a VM. With most vSphere 5.0 installations, the first bottleneck that you hit is memory. vSphere 4.1 supported a VM with as much as 255GB of memory, but vSphere 5.0 now supports a VM with as much as 1TB of memory! You can actually purchase a server now that can hold 2TB of memory, but it would probably take a month to boot. (Of course, it isn't the host that has the 2TB limit, but vSphere 5.0—so vSphere will see only 2TB even if your host has more than that.) Some of our clients have servers that require more than 255GB of memory, but none of them have servers that require more than 1TB of memory. You might not need 1TB of memory in a VM now, but it's good to know you have room for growth.

#3: As many as 32 vCPUs with Enterprise Plus. With vSphere 5.0 Enterprise Plus, you can have a VM with as many as 32 vCPUs. With the Essentials, Essentials Plus, Standard, and Enterprise editions of vSphere 5.0, you can configure a VM with as many as 8 vCPUs. In addition to vCPUs on a VM, you can also specify the number of cores in each vCPU. For VM applications that aren't SMP-aware, adding additional vCPUs to the VM doesn't usually improve the VM's performance. However, for VM applications that *are* SMP-aware (e.g., Exchange Server 2010), you can make a single vCPU more powerful by configuring it with multiple cores.

Be aware that vSphere 5.0 won't let you exceed the number of physical CPU cores on the host. If you have a two-socket, six-core CPU, then you can allocate 12 cores per VM, at most. When you configure the vCPUs on a VM, you'll see that as you increase the number of vCPUs, the number of cores decreases and vice versa. Using our previous example of the 12-core host running vSphere 5.0 Enterprise Plus, you can have a VM that has as many as 12 vCPUs with 1 core, 1 vCPU with 12 cores, or any valid combination within this range. The ability to specify both the number of vCPUs and cores within each vCPU gives vSphere 5.0 administrators an ability to fine-tune the performance of a VM that wasn't so easy to discover in vSphere 4.1.

#4: Support for VMFS5. With VMware VMFS3, you can have a maximum extent size of just less than 2TB. If you create an extent that is exactly 2TB, you might have difficulty seeing the partition on the ESX host. With VMFS3, you can have as many as 32 extents per storage group, so the largest storage group that you can configure is 64TB. However, we maintain a one-to-one extent-to-storage group relationship to simplify storage management. With VMFS3 you had to plan for the largest *.vmdk file or VM disk that you wanted to store on the storage group, and then format the storage group with the appropriate block size. Some ESX administrators always format their storage groups with an 8MB block size, to ensure that they can always have the largest *.vmdk file. The default block size with VMFS3 is 1MB, so the largest *.vmdk file you can have is 256GB. The relationship between VMFS3 block size and maximum *.vmdk file is shown in Table 2.

With VMFS3, you never want to create the maximum size *.vmdk that the storage group allows or you won't be able to take a snapshot of the VM. When a snapshot is

Table 2: VMFS3 Block Size vs. Maximum *.vmdk Size	
VMFS3 Block Size	Maximum *.vmdk Size
1MB (default)	256GB
2MB	512GB
4MB	1TB
8MB	2TB

created, it increases the base *.vmdk slightly, so if you're already at the maximum *.vmdk, then the snapshot creation will fail.

vSphere 5.0 supports VMFS5, which has significant improvements over VMFS3:

- **Maximum storage group size**—With VMFS5 partitions, the maximum extent size is 64TB.
- **Block size**—By default, VMFS5 partitions are formatted with a 1MB block size, but you can still have a maximum *.vmdk size of 2TB. In other words, you don't need to worry about the relationship between block size and maximum *.vmdk size anymore.

- VMFS5 sub-blocks—Sub-block size in VMFS5 has been reduced from 64KB to 8KB, so storage of smaller files takes up less space on the VMFS5 partition.
- File limit—File limit has increased from 30,720 files with VMFS3 to more than 130,000 files with VMFS5.

You can perform a nondestructive upgrade of an existing VMFS3 partition to VMFS5, but you will retain the existing block size of the VMFS3 partition, sacrifice the benefits of the smaller sub-blocks, and still have a file limit of roughly 30,000 files. The best practice is to create a new storage group that's formatted with VMFS5 and then use VMware Storage vMotion to migrate the VMs to the new storage group. Even though an in-place migration to VMFS5 is supported, it's a risky operation and can have potentially catastrophic results if something goes wrong.

#5: SSD storage groups. vSphere 5.0 can recognize a storage group that comprises solid state disks (SSDs) and can use it for memory swapping. As a general rule, we avoid using the memory overcommit feature in vSphere because it can hurt the performance of the VMs that are running on the host. But what if you need to run more VMs on an ESXi host on which the maximum amount of memory is already installed? SSD storage groups might be an option. Although they aren't as fast as native memory installed on the host, SSDs are still significantly faster than non-SSD storage. Enterprise SSDs are still expensive—around \$3,400 for a 200GB mainstream drive and \$7,000 for a 200GB enterprise performance drive. However, this might be a cost-effective solution when you're faced with replacing multiple ESXi hosts in a cluster. (Of course, there are drawbacks to this approach as well. Obviously, the best arrangement depends on your specific situation.) The other place we're using SSD drives on ESXi is for an application that requires fast disk performance. One of our clients has a legacy ERP application that runs only on Microsoft SQL Server 2000. SQL Server 2000 supports a maximum of 4GB of memory on

the server, 2GB of which is available for SQL Server. If the client could run a later version of SQL Server on the x64 platform, we would just load the server with a lot of memory and cache everything. Because this isn't an option and the client still needs the scalability, we'll place the SQL Server 2000 database on SSD storage to give the application significantly better performance.

Migrate Now for Improved ROI

There are numerous improvements in vSphere 5.0, and most deployments should be able to fit within the confines of the vSphere 5.0 vRAM entitlement model—although some vSphere users might need to purchase additional licenses. We suggest migrating to vSphere 5.0 today to get better ROI on your virtualization infrastructure investment. ■

InstantDoc ID 143333



Windows IT Pro

AUGUST SCHEDULE OF EVENTS
Web Seminars | In-Person Events | eLearning Events

DevPro **SharePoint Pro** **SQL Server PRO**

All Times Eastern

Web Seminars				
Tuesday, August 14, 1:00 PM Deployment Best Practices in Consideration of Windows 8	Wednesday, August 15, 12:00 PM Looking Ahead: Windows Server 2012 and Virtual Networking Environments with John Savill	Wednesday, August 22, 2:00 PM Web-enable your IBM i applications	Thursday, August 23, 12:00 PM Providing the Maximum Level of Malware Protection for Any Size Organization	Monday, August 27, 12:00 PM Idera SQL diagnostic manager gets an upgrade. You get the benefits.

E-Learning

Thursday, August 23, 11:00 AM
Preparing for SharePoint V15—Based on Two Years of Lessons Learned

Shared-Nothing VM Live Migration with Windows Server 2012 Hyper-V

Enjoy the benefits of consolidation—without dealing with downtime



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCITP: Enterprise Administrator for Windows Server 2008. He's a senior contributing editor for *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley, forthcoming 8/2012).

Email



Twitter



Website



In most aspects of computing, the goal is to consolidate, share resources, and function as a single unit. This consolidation is a key focus of virtualization. It was also a major priority for Windows Server 2008 R2 Hyper-V, which introduced the ability to share SAN-hosted NTFS volumes between all the nodes in a cluster. This approach gives all the volumes simultaneous access through the Cluster Shared Volumes (CSV) feature. Server 2008 R2 also introduced live migration, which lets you move a virtual machine (VM) between nodes in a cluster without downtime. Live migration does this by copying the memory and device state of the VM while it's running. The ability to move VMs with no downtime put Hyper-V on par with other hypervisors and gave organizations greater flexibility in placement and optimization of resources. For many organizations that use Hyper-V, live migration has become so important that one of the criteria for Windows Server 2012 was that no feature could be added if it would break the live migration capability.

New Capabilities

Server 2012 Hyper-V features numerous major new capabilities, including some that involve failover clustering. Two key changes relate to live migration within a failover cluster:

- Live migration supports multiple concurrent migrations between pairs of hosts in Server 2012, instead of just one.
- Failover clustering supports 64 nodes and 4,000 VMs in one Server 2012 failover cluster—a 400 percent increase over Server 2008 R2 failover clusters. This is not what I want to talk about in this article, though.

Server 2012 introduces a new type of live migration: a shared-nothing live migration, and I mean shared-*nothing*. No shared storage, no shared cluster membership—all you need is a Gigabit Ethernet connection between the Server 2012 Hyper-V hosts. With this network connection, you can move a VM between Hyper-V hosts, including moving the VM's virtual hard disks (VHDs), memory content, processor, and device state with no downtime to the VM. In the most extreme scenario, a VM running on a laptop with VHDs on the local hard disk can be moved to another laptop that's connected by a single Gigabit Ethernet network cable.

A word of caution: Do not think that shared-nothing live migration means that failover clustering is no longer necessary. Failover clustering provides a high availability solution, whereas shared-nothing live migration is a mobility solution that gives you new flexibility in the planned movement of VMs between Hyper-V hosts in your environment. As such, live migration can supplement failover clustering. Think of being able to move VMs into, out of, and between clusters and between standalone hosts without downtime. Any storage dependencies are removed with shared-nothing live migration.

Requirements for Shared-Nothing Live Migration

The requirements for enabling shared-nothing live migration are fairly simple:

- You need two (at a minimum) Server 2012 installations with the Hyper-V role enabled, or you need the free Microsoft Hyper-V Server 2012 OS.

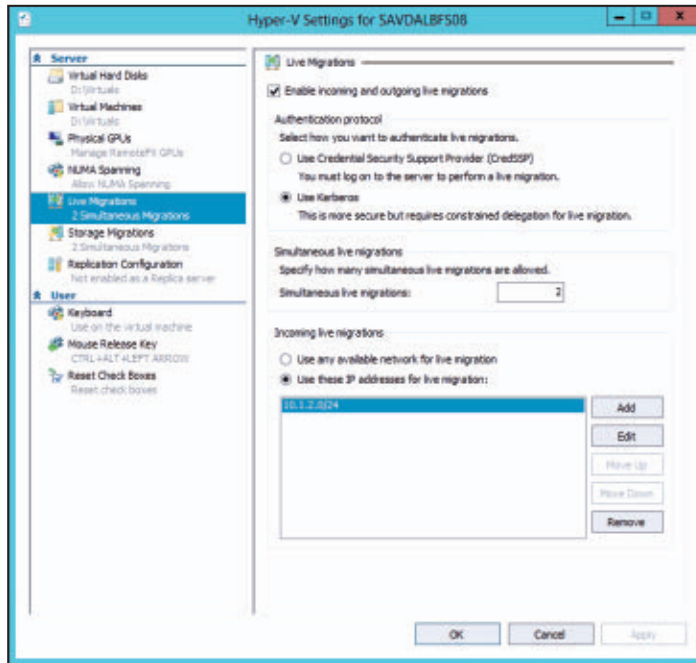
- Each server must have access to its own location to store VMs. This location can be local or SAN-attached storage or a Server Message Block (SMB) 3.0 share.
- Servers must have the same type or family of processor (i.e., Intel or AMD) if you're using the VM's Processor Compatibility feature.
- Servers must be part of the same Active Directory (AD) domain.
- Servers must be connected by at least a 1Gbps connection (a separate private network for live migration traffic is recommended but not necessary), over which the two servers can communicate. The network adapter that you use must have both the Client for Microsoft Networks and the File and Printer Sharing for Microsoft Networks options enabled, as these services are used for any storage migration.
- Each Hyper-V server should have the same virtual switches defined with the same name, to avoid errors and manual steps when performing the migration. If a virtual switch isn't defined on a target Hyper-V server that has the same name as the switch that's used in the VM configuration that's being migrated, then an error will be displayed and the administrator performing the migration will need to select which switch on the target Hyper-V server the VM's network adapter should connect to.
- VMs that are being migrated must not use pass-through storage.

Provided that you meet these requirements, the next step is to enable the Hyper-V hosts for incoming and outgoing live migrations.

Enabling the Hosts

To enable Hyper-V hosts for live migrations, select the *Enable incoming and outgoing live migrations* check box within the Hyper-V Settings for the host; these settings are available through Hyper-V Manager. Figure 1 shows the basic settings to enable live migration outside of a cluster environment.

In the simplest environments, selecting the *Enable incoming and outgoing live migrations* option, accepting the default Use Credential

**Figure 1**

The Hyper-V server options for live migration

Security Support Provider (CredSSP) setting for authentication, and using any available network for live migration should enable shared-nothing live migrations. Behind the scenes, a firewall exception for TCP port 6600 is enabled through the built-in exception Hyper-V (MIG-TCP-In). If you have a different local firewall on your servers or have firewalls between servers, then you must manually allow this port.

Note that in the Hyper-V Settings, you can also set the maximum number of simultaneous live migrations. Server 2012 removes the single simultaneous live migration limit between any two hosts and instead limits the number of simultaneous live migrations according to available network bandwidth, to give the optimal live migration experience. However, if you want to limit the number of simultaneous live migrations to a specific number, then set that limit in the *Simultaneous live migrations* field. You can also configure this setting by using the `MaximumVirtualMachineMigrations` parameter of the `Set-VMHost` Windows PowerShell cmdlet.

Although the default settings might work in simple environments or for a basic test, most environments will want to switch to Kerberos for authentication and will want to use a specific network for live migration traffic, which will include both a copy of the VM memory and its storage. Using Kerberos allows administrators to initiate live migrations remotely; using a specific network helps to manage network traffic and to ensure that the required network bandwidth is available for live migrations. Let's look at authentication first and why it's a challenge for live migration in a non-clustered environment.

Authentication for Live Migration

In a cluster environment in which all Hyper-V hosts are part of a failover cluster, all the Hyper-V hosts share a common cluster account. This account is used for communication between the hosts for authentication, simplifying (from an authentication perspective) operations such as migrations within a cluster. Outside of a cluster, each Hyper-V host has its own computer account, without a shared credential; when operations are performed, the user account of the user who is performing the action is typically used for authentication.

With a live migration, actions are taken on the source and target Hyper-V servers (and on file servers, if the VM is stored on an SMB share), both of which require the actions to be authenticated. If the administrator who is performing the live migration is logged on to the source or target Hyper-V server and initiates a shared-nothing live migration from the local Hyper-V Manager, then that administrator's credentials can be used both locally and to run commands on the target Hyper-V server. In this scenario, CredSSP works fine, allowing the administrator's credentials to be used on the remote server from the client—basically a single authentication hop from the local machine that is performing the action to a remote server.

However, the whole goal for Server 2012 (and management in general) is remote management and automation. Having to actually log on to the source or target Hyper-V server each time you require a

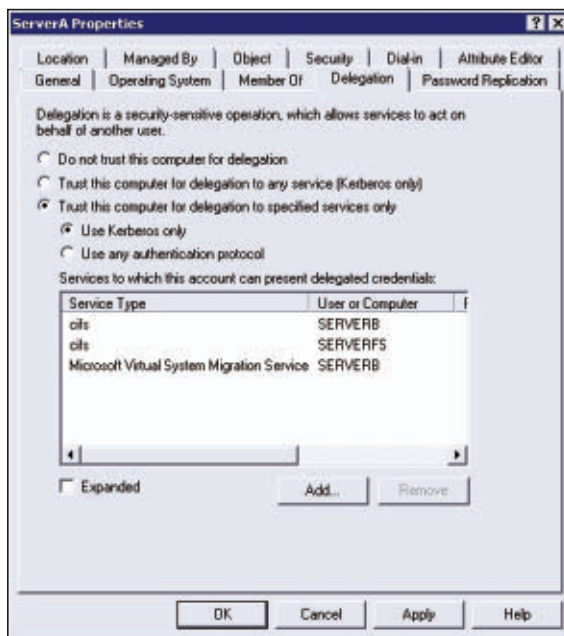
The whole goal for Windows Server 2012 (and management in general) is remote management and automation.

live migration outside of a cluster is a huge inconvenience for remote management. If a user was logged on to the local computer running Hyper-V Manager and tried to initiate a live migration between Hyper-V hosts A and B, that attempt would fail. The user's credentials would be used on Hyper-V host A (which is one hop from the client machine), but Hyper-V host A would be unable to use those credentials on Host B to complete the live migration. The problem is that CredSSP doesn't allow credentials to be passed to a system that is more than one hop away. This is where the option to use Kerberos enables full remote management: Kerberos supports constrained delegation of authentication. Therefore, when a user performs an action on a remote server, that remote server can use the user's credentials for authentication on a second remote server.

Does this mean that a server to which I connect remotely can just take my credentials and use them on another server without my knowledge? This is where the constrained part of *constrained delegation* comes into play, although you'll need to perform some setup before you can use Kerberos as the authentication protocol used for live migration. You need to configure delegation for each computer account that will be allowed to perform actions on another server on behalf of users. To configure this delegation, use the Active Directory Users and Computer management tool and the computer account properties of the server that will be allowed to delegate. As Figure 2 shows, the Delegation tab contains settings for the allowed level of delegation. For most computers, the configuration that this figure shows—allowing delegation only for specific services and only for the Kerberos protocol—is optimal. The only service that requires delegation is the Microsoft Virtual System Migration Service, which should be enabled for the target Hyper-V server. You *must* set authentication to *Use Kerberos only*. My two Hyper-V servers are SERVERA and SERVERB; the figure shows that I am modifying the delegation properties for SERVERA and configuring Kerberos delegation for the Microsoft Virtual System Migration Service to my other server, SERVERB. I'll repeat this configuration

Figure 2

The delegation settings to perform a remote live-migration initiation



on the SERVERB computer account, allowing it to delegate to SERVERA. Also note that I have delegation set for the Common Internet File System (CIFS) service, which is required later when VMs that are hosted on SMB file shares are migrated between hosts. After I've configured Kerberos delegation, live migration can be initiated between trusted hosts from any remote Hyper-V Manager instance.

Remember that all the hosts that participate in the live migration must have the same authentication configuration. Figure 3 and Figure 4 summarize the difference between CredSSP and Kerberos and the configurations that are required for each. Figure 3 illustrates the use of CredSSP, which requires the live migration to be initiated from one of the Hyper-V hosts. Figure 4 illustrates the use of Kerberos authentication and remote initiation, which requires the additional AD Kerberos constrained delegation. Although more work is involved in the use of Kerberos authentication, the additional flexibility makes the work worthwhile and definitely recommended. To configure the authentication type from PowerShell use the Set-VMHost cmdlet and set the VirtualMachineMigrationAuthenticationType to either CredSSP or Kerberos.

Network Settings

Authentication was the difficult part. Next, you must set the network to use incoming live migrations (i.e., the network on which the host will listen and accept live migrations). By default, live migration is




Figure 3
Migration using
CredSSP

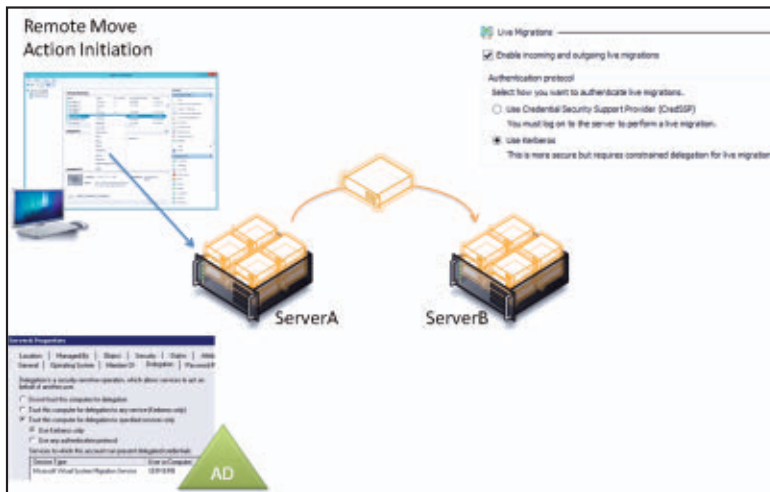


Figure 4
Migration using
Kerberos

accepted from any network. However, I recommend that you use a private, live migration-specific network whenever possible, to ensure that bandwidth is available and separate from other network traffic. You can add and order multiple networks: Simply enter the appropriate IP subnet, using the network prefix notation, also known as the Classless Inter-Domain Routing (CIDR) notation. For example, to specify my network adapter with IP address 10.1.2.1 and subnet 255.255.255.0, I'd use the notation 10.1.2.0/24. An alternative is to specify the full IP address with a subnet of 32 (e.g., 10.1.2.2/32), removing any ambiguity but requiring the configuration to be changed any time the IP address changes. Make sure that the source and target IP servers can communicate with each other by using the IP addresses that you specified for use with live migration, or the live migration

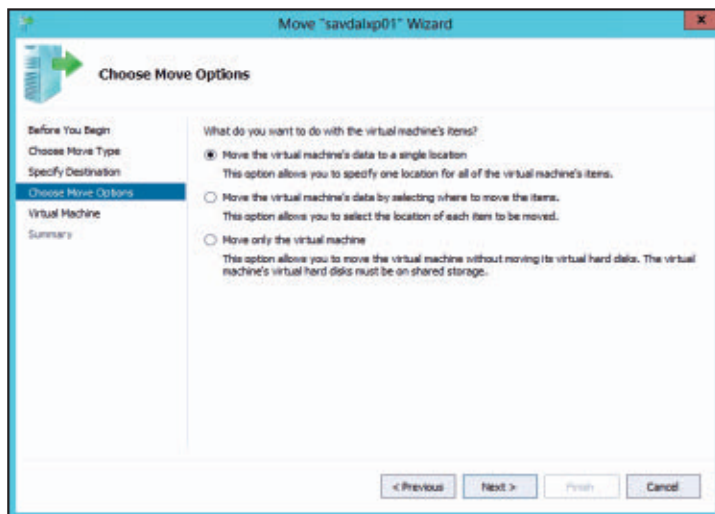
will fail. To configure these settings by using PowerShell, use the `Add-VMMigrationNetwork` and `Set-VMMigrationNetwork` cmdlets.

Using Shared-Nothing Live Migration

After you configure the Hyper-V hosts, the actual migration is simple. Select the Move action for a VM, then select the *Move the virtual machine* option as the move type. Enter the name of the destination Hyper-V server to which you want to move the VM, and finally choose how the VM's assets, such as VHDs, are moved to the destination. Figure 5 shows the final move option. Because we're focusing on the shared-nothing scenario, which means no shared storage and no use of SMB file shares, we need to select one of the first two options: *Move the virtual machine's data to a single location* or *Move the virtual machine's data by selecting*

Figure 5

Selecting options for the move operation



After you configure the Hyper-V hosts, the actual migration is simple.

where to move the items. The first option allows you to specify a single location on the target that will store the VM configuration, hard disks, and snapshots. The second option allows you to specify a location for each VM item in addition to selecting which items should be moved.

After you make your choice, select a folder on the destination server. The move operation will start; the time needed to finish is based on the size of the VHDs and the memory that you're moving

and the rate of change. However, the move will be completed without any downtime or loss of connectivity to the VM, as you can see in the accompanying video. You can also initiate the move by using the Move-VM PowerShell cmdlet.



Video

Using shared-nothing live migration

Troubleshooting Live Migration

The following steps should help you troubleshoot any hiccups that you might experience:

1. First, make sure that you have adhered to the requirements that I listed at the start of this article.
2. Check the Event Viewer (Applications and Services Logs > Microsoft > Windows > Hyper-V-VMMW > Admin) for detailed messages.
3. Make sure that the IP configuration between the source and target is correct. The servers must be able to communicate. Try pinging the target live-migration IP address from the source server.
4. Run the following PowerShell command in an elevated session, to show the IP addresses that are being used for a server and the order in which they're used:

```
gwmi -n root\virtualization\v2 Msvm_
VirtualSystemMigrationService | select
MigrationServiceListenerIPAddressList
```

5. Make sure that the Hyper-V (MIG-TCP-In) firewall exception is enabled on the target.
6. The target server must be resolvable by DNS. Try running Nslookup on the target server. Also run

```
ipconfig /registerdns
```

on the target server and

```
ipconfig /flushdns
```

on the source server.

7. On the source server, use the following command to flush the Address Resolution Protocol (ARP) cache:

```
command arp -d *
```

8. To test connectivity, try a remote Windows Management Instrumentation (WMI) command to the target (the WMI-In firewall exception must be enabled on the target); for example

```
gwmi -computer <DestinationComputerName> -n root\
virtualization\v2 Msvm_VirtualSystemMigrationService
```

9. Change the IP address that is used for live migration. For example, if you're using 10.1.2.0/24, try changing to the specific IP address 10.1.2.1/32. Also check any IPsec configurations or firewalls between the source and target. Check for multiple NICs

on the same subnet, which could cause problems; try disabling one if you find any.

10. Set authentication to CredSSP and initiate locally from a Hyper-V server. If this solves the issue, then the problem is the Kerberos delegation.

The most common problems that I've seen are a misconfiguration of Kerberos or the IP configuration. Failing to resolve the target server via DNS will also cause problems.

Closing Thoughts and Next Steps

Shared-nothing live migration is the most extreme type of zero-downtime migration. However, there are other types of zero-downtime migration, such as storing VMs on an SMB file share that both Hyper-V hosts can access. This approach transfers memory and device state over the network, without moving the storage. And there is still live migration within a failover cluster, which can use shared SAN-based storage through the CSV file system (CSVFS).

If you're moving a VM between failover clusters or into or out of a failover cluster from a standalone Hyper-V host, you'll need to remove the VM from the cluster before migrating the VM. The good news is that with Server 2012, you can add and remove a VM from a failover cluster without needing to stop the VM—meaning no downtime to the VM, even when migration with a failover cluster is involved.

Of course, even in a shared-nothing scenario, there is still a shared physical network fabric and a dependence on the VM IP configuration during the move. This is where another Server 2012 feature, Network Virtualization, can open up a world in which VMs can be moved between any hosts in different locations without changing the networking configuration of the VM OS. ■

InstantDoc ID 143168

Administrative Reporting with PowerShell

Use custom objects to get a grip on data



Max Trinidad

is a SQL Server developer and has been a PowerShell MVP since 2008. He has worked with legacy systems, network infrastructures, and databases since 1979.

Email



Finally, IT professionals are realizing how important Windows PowerShell has become as part of their IT tools. This new technology has so many flavors that it takes time to learn them all. But thanks to our many Internet PowerShell communities, the learning curve can be reduced dramatically. Resources can be found in books, videos, conferences, and blogs all over the Internet. One more thing: This information isn't limited to IT admins. Everyone can learn and use PowerShell.

Now, we all know that PowerShell is an excellent tool for automating repetitive tasks. In this article, I'm going to show how to use PowerShell to create an administrative reporting task. At the same time, you'll learn some PowerShell during this article. But first, you need to meet some requirements:

- You have PowerShell 2.0 installed, with Integrated Scripting Environment (ISE) feature enabled on the server.
- You can use the PowerShell ISE Editor.
- You need Microsoft Office installed so that you can use Microsoft Excel to open comma-separated value (CSV) files.

PowerShell Version 2.0 is already included in Windows 7 and Windows Server 2008 R2. For earlier, legacy OSs such as Windows Vista Service Pack 1 (SP1), Windows XP SP3, Windows Server 2008 SP1, and Windows Server 2003 SP2, you need to download the [Windows Management Framework 2.0 RTM](#) to get PowerShell 2.0.

Getting Started

To start building your script solution, PowerShell provides you with a free editor: PowerShell ISE. Choose Start, All Programs, Accessories, Windows PowerShell ISE, as Figure 1 shows.

However, before you start building and executing scripts, you must understand a bit about the PowerShell execution policy. By default, this policy is set to Restricted, meaning that PowerShell will not execute scripts. The most important thing to understand is that the PowerShell execution policy is meant only to prevent you from harming the computer. The policy isn't meant to stop hackers or viruses. In most cases, setting the script execution policy to RemoteSigned is enough to get you working with PowerShell scripts:

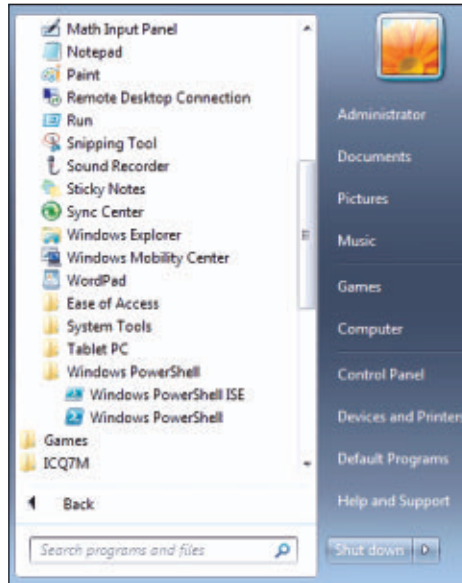


Figure 1
Opening
PowerShell ISE

```
Set-ExecutionPolicy RemoteSigned
```

Be aware that on Windows Server 2008 and later or Windows Vista or later, you must execute this command from an elevated session.

If you want to learn more about setting the execution policy, use the following command:

```
Help About_Execution_Policy -Full
```

PowerShell is loaded with Help information, all accessible at your fingertips. Here's another command, which displays a list of available Help content, as Figure 2 shows:

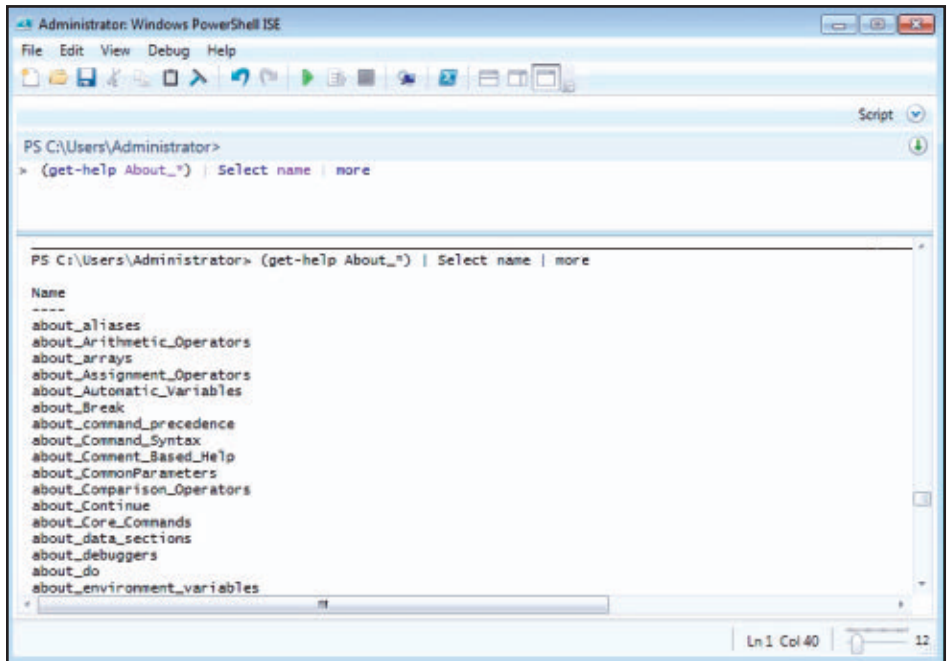
```
(get-help About_*) | Select name | More
```

All these commands are known as PowerShell oneliners.

Now that you have the basics of getting started, let's look at a practical scenario. Suppose our IT manager is performing an audit and requests an administrative report to gather the OS serial number and system IP address on three specific systems: Server1, Desktop1, and Server2. We'll need to follow these steps:

1. Input—Store the computer names.
2. Process—Use Get-WMIobject to gather information.
3. Output—Use Export-CSV to create an Excel CSV file.
4. Automation—Put it all together.

Figure 2
PowerShell Help
content



Input: Storing Computer Names

First, we need to store our three selected computer names in a PowerShell variable:

```
$ComputerList = @("Server1","Desktop1","Server2");
$ComputerList.gettype()
```

As you can see, the PowerShell variable name starts with a dollar sign (\$). This variable is a collection of Microsoft .NET Framework objects. Saving the series of computer name values, enclosed in quotation marks (") and separated by commas, is the simplest way to create a list string array. To check the type of variable that you created in PowerShell, you can add the .NET .gettype() method to the end of the variable.

By the way, if the list of computers comes in the form of a text file, then you can use the following command to create the \$ComputerList variable:

```
$ComputerList = Get-content c:\temp\ComputerList.txt
```

This command automatically creates a string array type PObject, and you can easily maintain the text file on a local or network drive.

Process: Using Get-WMIObject to Gather Information

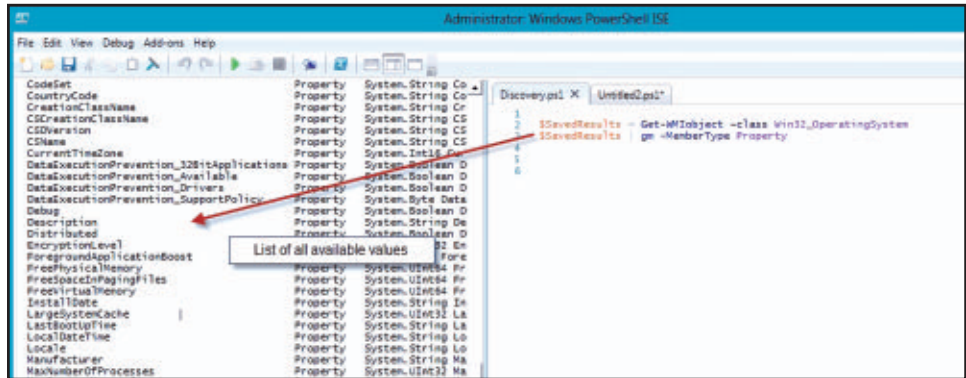
Here's where PowerShell excels in getting information. We are going to use the Get-WMIObject command. (Again, you can use the Help command to read about this particular command.) Based on our manager request, we have identified that we need to use two WMI classes to produce our report, which we'll call Win32_OperatingSystem, and one WMI query to retrieve the machine IP addresses.

It's a good idea to familiarize yourself with your command by writing a oneliner. Let's run the Get-WMIObject command against our local machine, creating a PObject variable. By saving the results to a variable, we can use the Get-Member command to display all the values that are saved in this .NET object, as Figure 3 shows.

```
$SavedAuditOS = Get-WMIObject -Class Win32_OperatingSystem
$SavedAuditOS | GM -MemberType Property
```

Figure 3

Displaying values in the .NET object



We use the pipe (|) to pass our variable results to another command, in this case GM. GM is an alias of the Get-Member command, which helps us to discover and expose all the fields (or properties) that are stored in our recently created variable.

From the `$SavedAuditOS` variable, we select the properties: `CSName`, `SerialNumber`, and `Name`. Here's our oneliner command to display the selected properties:

```
$SavedAuditOS | Select-Object CSname, SerialNumber, Name | FT
-AutoSize
```

We're now passing the variable to two other commands: Select-Object, and FT (alias for Format-Table). The Select-Object command helps you to display the value of the PSObject properties. The Format-Table command helps you to display the results in a table-formatted view and, with the help of the -AutoSize parameter, eliminates the additional spaces between the displayed columns.

Now, let us proceed in getting the IP address information. We're going to take a different approach with our next Get-WMIObject command by including the WMI -Query parameter:

```
$SavedAuditIPAddr = Get-WmiObject `
    -query "SELECT * FROM Win32_PingStatus WHERE
```

```
Address='$MyMachineName' "`
| select StatusCode, IPv4Address | ft -auto;
```

This block of code is still considered a oneliner, even though it uses the backtick operator to split and make the line more readable. Also, the semicolon (;) can be used to set the end of the line. In the WMI -Query string, make sure to change the variable following Address = to include the computer name.

Note that we're using Win32_PingStatus to gather information; we aren't using Test-Connection because it returns data only if you can ping the computer. In addition, you could shorten your code a bit:

```
Get-WmiObject Win32_PingStatus -filter
"Address=''$MyMachineName''"
```

Output: Using Export-CSV to Create an Excel CSV File

Here's the last piece of the puzzle, which will allow us to create an Excel CSV file. This will take only one PowerShell command, using the previously created \$SavedAuditOS variable:

```
$SavedAuditOS | Export-CSV -Path C:\temp\TestReport.csv
-NoTypeInfo;
ii C:\temp\TestReport.csv;
```

This is another oneliner, in which we include the ii (an alias for Invoke-Item) command. This command allows us to open a file with the associated installed application, in this case Excel. In most cases, Excel will be the default application for viewing a CSV file. Be aware, if you don't have a CSV application installed, you'll need to use Microsoft Notepad instead of Invoke-Item. The -NoTypeInfo parameter excludes extra information about the data that was added to our exported CSV file.

Automation: Put It All together

We've gone through the steps to basic understanding of how our commands work. But there is one more puzzle piece. We have three variables: one holding the list computer, two others with WMI computer information. Here's where the PowerShell magic starts.

We need to consolidate all computer information into a single variable named `$AuditOSInfo`. This variable is a PowerShell hash table that contains the results of both WMI processes: the `Win32_OperatingSystem` and the `Win32_PingStatus` query. Then, using the `ForEach` statement, we'll go through each computer name, building the hash table with all the information.

Inside the `ForEach` command, we'll create two variables:

- `$SavedAuditOS`—collects the `Win32_OperatingSystem` information from the selected Computer
- `$SavedAuditIpAddr`—collects the WMI query `IPAddress` information from the selected computer

Inside the `ForEach` code, we build a hash table object `$MyPSObject` to consolidate the information for each computer and add it our variable `$AuditOSInfo`. Finally, after collecting all our information, we use the `Export-CSV` command to create our output CSV file. Listing 1 shows the full script. The script will generate the sample administrative report that Figure 4 shows.

About Our Custom Hash Table Object

This is the heart of our process. Here is where we consolidate our data into one custom object, known as the hash table. We can customize our results, merging data from different sources, and changing the label of our data value when necessary. This all happens under the `-Property` parameter when we create our new `PSObject`, as the code in Listing 2 shows.

Note that I'm using an advanced technique when creating my `OSName` property to get the OS description value stored in the

Listing 1: Code to Create Output CSV File


[Download the code](#)

```
#####
## AdminOSReport.ps1
## 12/04/2011
#####

## - Create variable with list of ComputerNames:
$ComputerList = "WIN8Server1","WIN764SQL01","W764SQL02Merged";

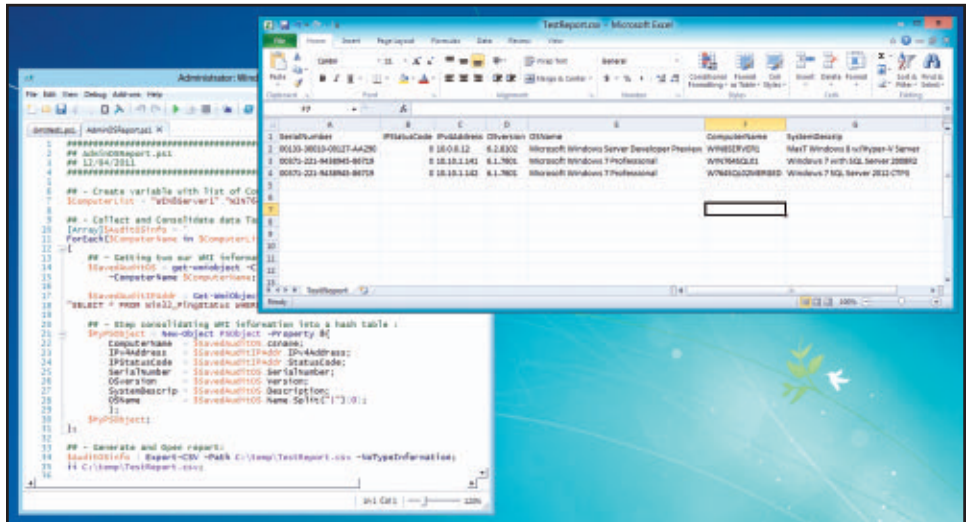
## - Collect and Consolidate data Task:
[Array]$AuditOSInfo = `
ForEach($ComputerName in $ComputerList)
{
    ## - Getting to our WMI information:
    $SavedAuditOS = get-wmiobject -Class Win32_OperatingSystem `
        -ComputerName $ComputerName;

    $SavedAuditIPAddr = Get-WmiObject -query `
"SELECT * FROM Win32_PingStatus WHERE Address =
'$ComputerName'";

    ## - Step consolidating WMI information into a hash table :
    $MyPSObject = New-Object PSObject -Property @{
        ComputerName    = $SavedAuditOS.csname;
        IPv4Address      = $SavedAuditIPAddr.IPv4Address;
        IPStatusCode     = $SavedAuditIPAddr.StatusCode;
        SerialNumber     = $SavedAuditOS.SerialNumber;
        OSversion        = $SavedAuditOS.Version;
        SystemDescrip    = $SavedAuditOS.Description;
        OSName           = $SavedAuditOS.Name.Split("|")[0];
    };
    $MyPSObject;
};

## - Generate and Open report:
$AuditOSInfo | Export-CSV -Path C:\temp\TestReport.csv
-NoTypeInfo;
ii C:\temp\TestReport.csv;
```

Figure 4
Sample administrative report



Listing 2: Code to Create Custom Hash Table Object

```
## - Step consolidating WMI information into a hash table :
$MyPSObject = New-Object PSObject -Property @{
    ComputerName = $SavedAuditOS.csname;
    IPv4Address = $SavedAuditIPAddr.IPv4Address;
    IPStatusCode = $SavedAuditIPAddr.StatusCode;
    SerialNumber = $SavedAuditOS.SerialNumber;
    OSversion = $SavedAuditOS.Version;
    SystemDescrip = $SavedAuditOS.Description;
    OSName = $SavedAuditOS.Name.Split("|")[0];
};
```

\$SavedAuditOS variable Name property using the String Split() method. This property contains three values, separated by a pipe (|), but we want only one, which is identified as element 0 in our code:

```
OSName = $SavedAuditOS.Name.Split("|")[0];
```

One of the three values will be stored in our custom hash table's OSName property. Taking this multi-value, pipeline-delimited string,

we use the `Split("|")` method followed by the square bracket `[0]` to tell PowerShell to return the first element of the string.

Here's an example of a multi-value string containing three delimited elements:

```
$Str_name = "Microsoft Windows 2008R2|C:\Windows\Device\
    Harddisk0\Partition2"
Element #  -> [0]  -> [1]  -> [2]
```

Using the example line `$Str_Name.Split("|")[0]`, we grab only the first element, selecting the string value *Microsoft Windows 2008 R2*. As you can see, this uses the .NET string `Split()` method in the scenario when you need to extract a piece of the data with a delimited character.

Get Creative

For the most part, this is not a sophisticated script and is meant to be executed from a desktop. There's no reason you need to be on the server to use PowerShell. We can still add more logic or code to add more fields. You can even make the script capable of sending an email message with the CSV file attached, or you can create a scheduled task to run the script at a set date and time. There are many ways to use this script, and you can be very creative.

On a side note, when creating your custom hash tables, you will notice that your .NET object properties aren't in the same order as was initially defined. This is normal in version 2.0. Just use either the `Select-Object` or `Format-*` command and type the properties in the order that you want them to be displayed.

Just remember, PowerShell has become an essential tool for administrators and can no longer be ignored. Take the time and get to know it. ■

InstantDoc ID 143019

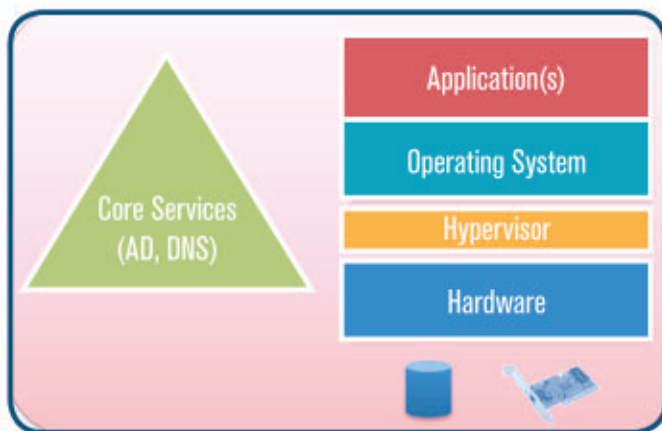
Why Monitoring the OS Only Is Not Monitoring At All

Many organizations have no monitoring implemented. The “alerting” solution for the organization is when the users—or even worse, the customers—complain that a service is not available, at which point troubleshooting and remediation begins. The result is prolonged service unavailability and loss of reputation and money. Other organizations realize the dangers of having no monitoring solution and implement only host ping monitoring. However, this is really not even half the picture; it covers only a very small set of failure conditions. A true monitoring solution needs to provide proactive monitoring

of all elements of a service, allowing resolution before a failure occurs where possible to avoid any impact to service availability. And the monitoring solution should provide detailed information and guidance to aid in resolution.

When choosing and implementing a monitoring solution it’s critical to consider all aspects of the service and the service’s dependencies. The monitoring solution should inform the organization of not just failures but conditions that could lead to a failure, misconfigurations, and performance degradation situations. This level of complete monitoring requires insight and intelligence at all levels of the service,

Complete Service Monitoring



Standard and custom application monitoring
Synthetic transactions and user experience evaluation
OS health and performance
Automatic discovery of applications
Hypervisor awareness
Early fault detection
Integration with hardware monitoring capabilities
Integration with storage and networking components

ABOUT SOLARWINDS SERVER & APPLICATION MONITOR

Solarwindows Server & Application Monitor is an industry leading monitoring solution for heterogeneous environments providing a single interface for management of all major operating systems, hypervisors, hardware platforms and applications in addition to featuring powerful extensibility. You can find detailed information at www.solarwinds.com/server-application-monitor.aspx, in addition to a 30-day free trial and an online test drive of the solution.

including the hardware, operating system, and the application. The complete service should be modeled in the monitoring solution to give an easy view of whether or not a complete service is healthy based on all the component hardware, operating systems, and applications that make up the service.

Application monitoring includes services such as databases, mail systems, web services, collaboration solutions, and foundational services such as Active Directory and DNS. A monitoring solution should have built-in capabilities for monitoring predominant applications and it's critical to not just check whether a process or service is running but also whether it's responding to user requests and responding within an acceptable amount of time which can be achieved through synthetic transactions. Many organizations also have applications that have been developed in-house so monitoring solutions should be customizable to perform monitoring of custom applications.

Virtualization of server operating systems is predominant in many organizations, which means monitoring of the hypervisor health is critical—so build such monitoring into your solution along with the monitoring of the operating systems that host the applications. Very few organizations use a single operating system, such as only Windows or only Linux, and the goal for monitoring should be a single pane of glass. Choose a solution that has the ability to natively monitor all the server operating systems used in your environment and, ideally, one that automatically detects the applications installed as the operating system is added to the monitoring solution. Ease of use and configuration should be a top priority.

Finally, consider the underlying hardware that allows operating systems and applications to run. While hardware has become more reliable and fault tolerant it has become more complex. Hardware monitoring provides the final level of insight to offer a complete monitoring solution. Some major hardware vendors have their own proprietary monitoring components, which means integration will be a key requirement to maintain the single monitoring pane of glass.

Getting the right monitoring solution in place will require some work and planning, but by selecting a solution that can monitor all the hardware, systems, and applications in your environment the process will be far simpler and, more importantly, far more powerful than a multitude of point monitoring products or no monitoring at all. ●

ALL THE POWER AT A PRICE *ANYONE* CAN AFFORD!



SolarWinds Server & Application Monitor

Powerful, agentless monitoring for less than yearly maintenance on most traditional solutions. **Starting at \$2,995!**

 **DOWNLOAD FREE 30-DAY TRIAL**

TEST DRIVE LIVE DEMO »



WINNER!

Best of Microsoft TechEd 2012
in Systems Management and Operations

solarwinds
Unexpected Simplicity™

Understanding App Controller 2012

It's all about the apps and services

Microsoft System Center App Controller is a new member of the System Center family of products. Although other products in this suite can be implemented independently of one another (with the ability to integrate, of course), App Controller is highly dependent on System Center Virtual Machine Manager (VMM) or Windows Azure. In case you aren't familiar with App Controller's purpose, let me make a brief introduction.

App Controller is a product for managing applications and services that are deployed in private or public cloud infrastructures, mostly from the application owner's perspective. It provides a unified self-service experience that lets you configure, deploy, and manage virtual machines (VMs) and services. Some people mistakenly think that App Controller is simply the replacement for the VMM Self-Service Portal. Although App Controller does indeed serve this function, and in some way can replace the Self-Service Portal, its focus is different. VMM Self-Service portal was used primarily for creating and managing VMs, based on predefined templates; App Controller also focuses on services and applications. App Controller lets users focus on what is deployed in the VM, rather than being limited to the VM itself.

To understand this concept, you need to be familiar with VMM 2012. Although this article is not about VMM, I must mention some important things so you can get the full picture. VMM 2012 has significantly changed from VMM 2008 R2. VMM 2012 still manages and deploys hosts and VMs, but its main focus is on private clouds and service templates. The end result is that an administrator or end user



Damir Dizdarevic

is manager of the Learning Center at Logosoft in Sarajevo, Bosnia and Herzegovina. He's an MVP for Windows Server Infrastructure Management, and an MCSE, MCTS, MCITP, and MCT. He works as a system designer for enterprise environments, and he's also an author of several Microsoft Official Courses.



Email



Twitter



LinkedIn



Facebook



Blog
(in Bosnian language)

App Controller lets users focus on what is deployed in the VM, rather than being limited to the VM itself.

can deploy a service or application to a private cloud even without knowing exactly what lies beneath it.

I mentioned earlier that you can use App Controller to connect to both private and public clouds. Connecting to a private cloud means establishing a connection to a VMM 2012 Management Server. However, you can also add a Windows Azure subscription to App Controller.

Target users for App Controller are not administrators, although some admin tasks can be performed through the App Controller console. App Controller is intended to be used by application or service owners: the people that deploy and manage an application or service. (Don't confuse these folks with the end users that actually use a service or application. End users should not be doing anything with App Controller.) An owner might be an administrator, or an owner might be a developer that needs a platform to test an application. The key point is self-servicing: App Controller enables application owners to deploy new instances of a service or application without requiring them to deal with jobs such as creating VMs, Virtual Hard Disks (VHDs), or networks or installing OSs. To achieve that level of automation, administrators should do a lot of work in VMM.

App Controller can't create or manage building blocks for VMs or services. Nor can it be used to create new objects from scratch (except for service instances). Anything you work with in App Controller must first be prepared in VMM. That means creating VM templates, guest OS profiles, hardware profiles, application profiles and packages, and logical networks, as well as providing Sysprepped .vhd files, ISO images, and private cloud objects. To deploy services through App Controller, a VMM administrator must create a service template and deployment configuration. Self-service user roles also should be created in VMM and associated with one or more private clouds and quotas.

App Controller doesn't have its own security infrastructure: It relies completely on security settings in VMM, so available options for a

user in App Controller depend directly on the rights and permissions that are assigned to the user in VMM. Authentication is performed by using a web-based form, but you can opt to use Windows Authentication in Microsoft IIS to achieve single sign-on (SSO).

Installation and Initial Configuration

App Controller is a lightweight product. The installation image is only 30MB and runs as a Microsoft Silverlight web application. You can install it on any domain-joined Windows Server 2008 R2 machine with the Web Server role installed. You'll also need Microsoft .NET Framework 4.0, as well as Microsoft SQL Server for creating the App Controller database. (You can safely reuse the same SQL Server instance that you used for VMM installation.) Optionally, you can also install the PowerShell Module for App Controller if you want to manage it in that way.

App Controller can be installed on the same machine as VMM Management Server, which is probably one of the most common scenarios. However, if you decide to deploy App Controller as a stand-alone server, you'll need to install VMM Management Console as a prerequisite. To access App Controller, you need a web browser with the Silverlight client installed. For more information about these requirements, see the [Microsoft System Center App Controller System Requirements](#) page.

Installation of App Controller is a fairly simple procedure. After selecting the installation path, you should configure an account for the App Controller service to use. You can use the Network Service account, which is the default, or you can specify a domain account (consider using Managed Service accounts) that you create for this purpose. You can also change the port that App Controller uses for its internal communication (the default is 18622). Because App Controller is a web-based application, HTTP traffic is not allowed and you should also select an SSL certificate. App Controller setup can create a self-signed certificate, or you can use another certificate, such as

one that your private Certification Authority (CA) issues. Finally, you should select the SQL Server instance that will be used and choose a database name. After a minute or two, you're done.

The first thing to do after installing App Controller is to log on to it and connect to the private or public cloud. Open the App Controller website, log on as an administrator, and enter the Overview page. Click the option *Connect to a Virtual Machine Manager server and clouds* and a new window will open to let you add new VMM connections. You can add as many as five VMM connections. You should type the Fully Qualified Domain Name (FQDN) of the VMM server, as well as a connection name (which can be whatever you want). Port value should be left at the default value (8100) if you didn't change it on the VMM side. There is also an option to automatically import SSL certificates from the VMM server, which is needed to import files and templates from VMM to App Controller. If you don't have any security issue with this option, then you can leave it enabled.

Besides connecting App Controller to VMM, you can also make a connection to Windows Azure if you subscribe to that service. This step is optional but recommended if you use Windows Azure. Adding Windows Azure is a bit tricky: You should also provide a connection name and type to your Subscriber ID. However, because you can't automatically import certificates from Windows Azure, you should first go to Windows Azure Management Portal, click the Hosted Services, Storage Accounts & CDN tab, and then select Management Certificate in the navigation pane. There, you should add a certificate to your Windows Azure subscription. Windows Azure allows you to create your own management certificates, either self-signed or by using their preferred CA. Whatever you choose, export the .pfx file and then select it in App Controller as a management certificate when configuring the connection to Azure. App Controller stores this certificate in the App Controller database. Because the certificate contains the private key, you must provide the password so that App Controller can use the private key. (For a step-by-step discussion of this procedure,

see the article “[Q: How do I create a certificate to enable System Center App Controller to manage Windows Azure?](#)”)

Certificates are used to set up the trust between the Windows Azure management API and App Controller. This trust allows App Controller to call on the Windows Azure API when tasks such as deploying services or changing configuration properties are performed in the App Controller console. The management certificate (.cer file) contains only the public key, which is kept in Windows Azure for accessing the API. By giving Windows Azure the public key and keeping the private key local, the authentication can be completed.

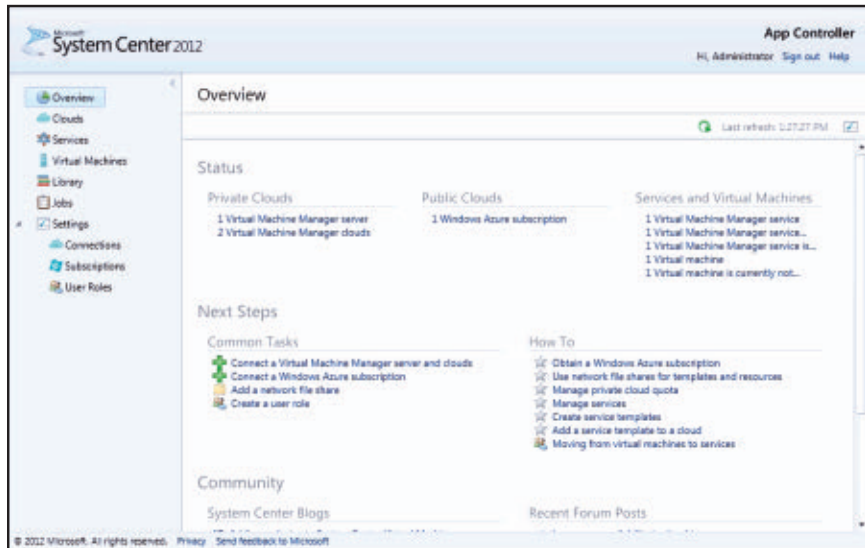
Anything you work with in App Controller must first be prepared in VMM.

Browsing the App Controller Console

After you finish the initial configuration of App Controller, you can start to use the App Controller console. If you log on to App Controller as an administrator, you can perform all available tasks. But as I said at the beginning, there is not much point in using App Controller as an administrator, except for during initial configuration. Administrators can perform all App Controller tasks (and much more) through the VMM console. To enable other users (i.e., application and service owners) to use the App Controller console, you should first add them to the Self Service user role in VMM. You should then define the scope and available resources (e.g., private clouds, VM templates, storage) for that user role. Also, define quotas for self-service users so that you can keep resource usage under control. (If you don't properly complete these steps, you can easily run into a situation in which App Controller users quickly fill all available resources on your VM hosts—probably a scenario that you don't want.) When a self-service user logs on to App Controller, that user will see only the resources and actions that you configured inside VMM.

The first thing that you see when you log on to App Controller is the Overview pane, which Figure 1 shows. It gives you status information about available private and public clouds, as well as about services and VMs to which you have access. You can also find quick

Figure 1
The App Controller
Overview pane



links to deploy new services or VMs. And if an Internet connection is available, the Overview pane shows the most recent System Center–related blog posts, as well as any forum posts from the VMM Forum on Microsoft TechNet. This information can be useful, so it’s great that Microsoft included it. You can also find how-to links for some common tasks, which is great for new users. If you log on as an administrator, you can make new VMM or Azure connections from this pane, as well as create new user roles and add network file shares.

The second pane is the Clouds pane, which Figure 2 shows. In this pane, self-service users can see the private and public clouds to which they have access. Users can start deployment of a new VM or service into the cloud and can manage Run As accounts. If a user chooses to deploy a new VM or service from this point, then that user is presented with a new deployment diagram, as Figure 3 shows, and can choose a VM or service template and start deployment. From the Clouds pane, an administrator can see all available private and public clouds.

On the Services pane, which Figure 4 shows, users can see the services and service instances that they’ve deployed. From here, users

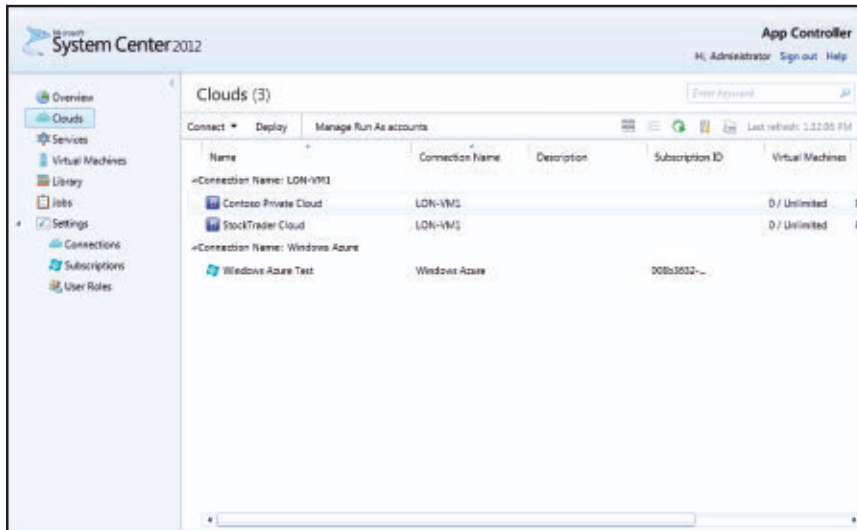


Figure 2
The App Controller Clouds pane

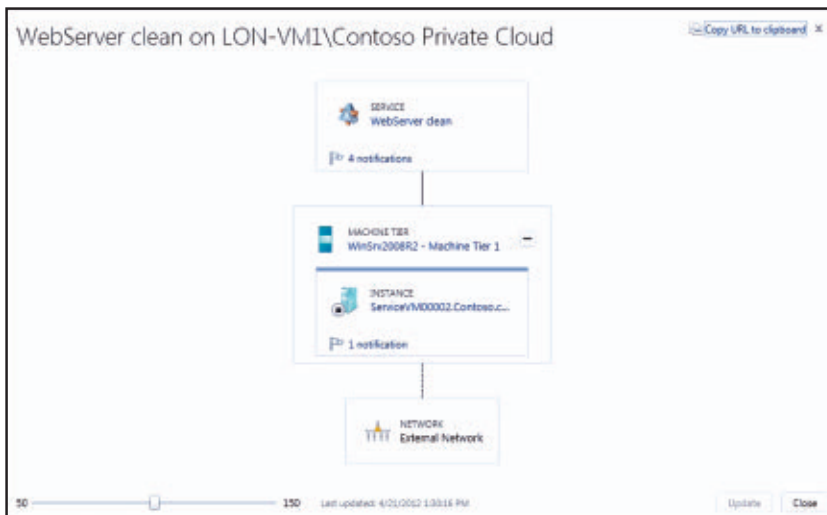
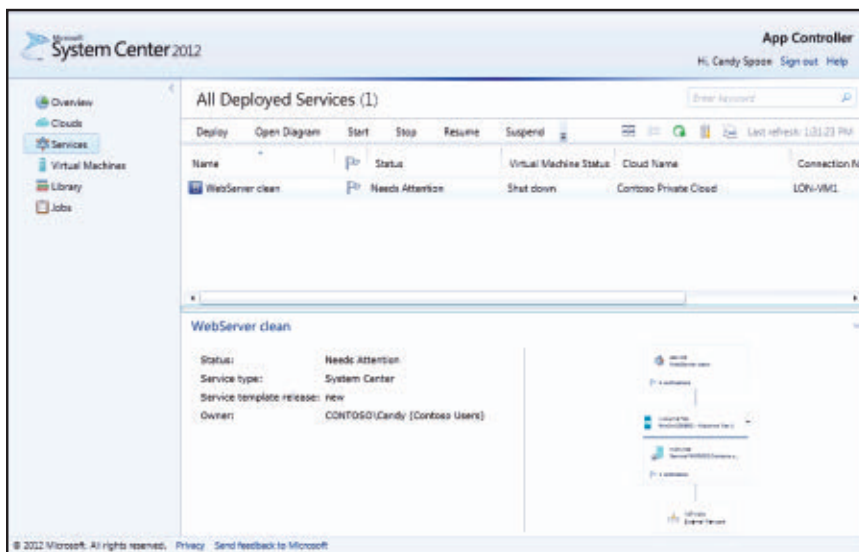


Figure 3
A new deployment diagram

can also deploy a new service; open a diagram of an existing service; and start, stop, resume, suspend, or shut down existing services. Be aware that these actions apply to a *service*, not to a specific VM. In fact, when you perform an action on a service, one or more VMs are indirectly affected by that action. For example, if you decide to shut down a service, all the VMs that are associated to that service will

Figure 4
The App Controller
Services pane



shut down. Actions that are available to a user here directly depend on the allowed actions in the VMM Self-Service user role settings. From this pane, you can also initiate a service upgrade (if available), fix errors, and delete service instances. Administrators can see all services that users have deployed, with a full set of available actions.

The Virtual Machines tab shows the existing VMs. For a self-service user, only VMs that are in the scope of the user's role are shown, as well as VMs that are used for any services that user deployed. From this pane, users can perform similar actions as on the Services tab, but on a VM basis. Besides regular actions, such as start, turn off, shut-down, and so on, you can also mount an ISO image into a VM and connect to a VM desktop by using RDP, as well as see VM properties.

The Library pane, which Figure 5 shows, gives self-service users access to VMM Library resources. Users can use this pane to see templates that are shared with them, as well as to access shared folders. Regular self-service users can't make any changes here but can use available resources to deploy new VMs or services and to access some shared files. Administrators can use the Library pane to create new shares that are available to App Controller and to perform some

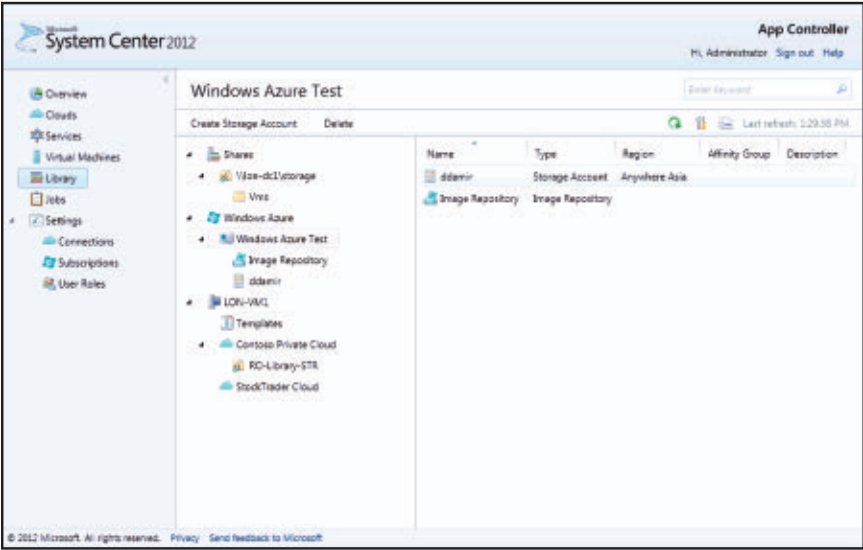


Figure 5
The App Controller Library pane

modifications on existing resources. The Library pane also shows available resources in the public cloud (Azure) and lets administrators create new storage accounts and browse image repositories.

The final pane in the App Controller console is the Jobs pane, which Figure 6 shows. This pane has primarily the same functionality

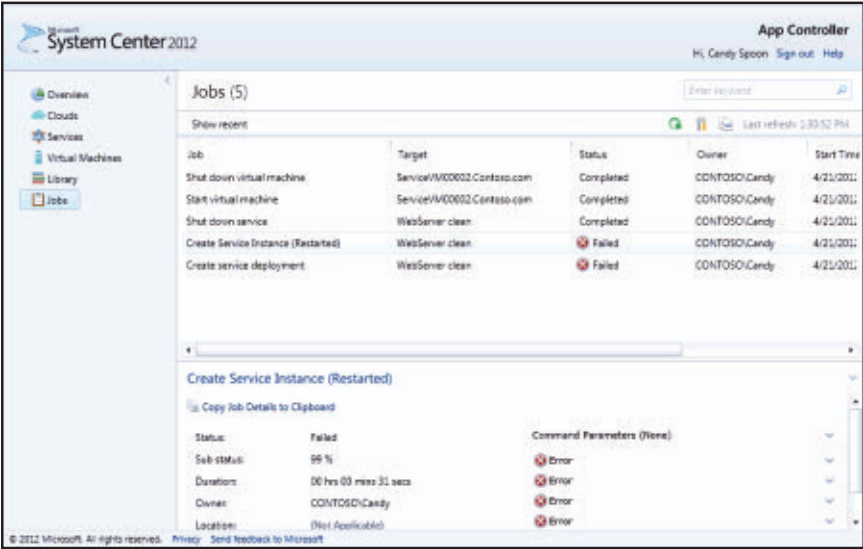


Figure 6
The App Controller Jobs pane

as the Jobs pane in VMM. Each self-service user can browse through jobs that the user initiated; administrators can see all jobs that have been performed through App Controller. You can also see details for each job.

One more pane is available in the App Controller console, but only for administrators. The Settings pane lets administrators manage connections to private clouds and connections and subscriptions for public clouds. Also, administrators can use the Settings pane to create User roles. However, unlike other options, these roles aren't connected to VMM. User roles that are created here are not synced to VMM, nor can you see VMM User roles on this pane. The purpose of user roles in App Controller is to provide users with access to the public cloud infrastructure. Because self-service users can't add connections to a public cloud or directly use connections that administrators make, this pane is the only way to provide those users with access to public cloud resources. However, the use of the same terminology for both VMM user roles and App Controller user roles can be a bit confusing.

Ready When You Are

App Controller is no doubt a useful product. It isn't for everyone, but if you want to enable self-servicing and have users with an appropriate level of knowledge to use App Controller, it can be beneficial. (Companies with such scenarios will probably also want to consider System Center Service Manager as an additional layer for managing, requesting, and approving new virtual environments.) The transition to the concept of services and applications has started, but a full conversion won't happen quickly. From that perspective, App Controller is definitely a next-generation product. However, you can still use it now as a replacement for the VMM Self-Service Portal, even if you don't yet want to make the transition to services. ■

InstantDoc ID 142925

SharePoint Security 101

Protect and preserve content with these security basics

For most organizations, a healthy security stance is essential to protecting and preserving Microsoft SharePoint content.

Depending on your industry, you might even be bound by regulations that require you to secure and audit access to some content. Certainly, security is a broad topic that involves not only SharePoint farm administrators, but server, network, and database admins as well. Through delegation to site collection admins and site owners in the business units, security also involves day-to-day end users. To effectively enforce business-driven information-management policies, collaboration with the business is also necessary.

In response to the vast number of interested security stakeholders, SharePoint 2010 offers a flexible security model that supports multiple authentication types and multiple levels of permission authorization. But this flexibility has a downside: Without a clear, published strategy for how permission management should be performed in your organization, security chaos can make day-to-day usage quite painful. With all these variables, it's no wonder that security can become a complex equation.

Rest assured that armed with a clear understanding of your security requirements, some technical guidance, and bit of common sense, you *can* build a security model that is both effective and sustainable. In this three-part series, we aim to provide you with this technical guidance—and we'll throw in some best practices for good measure. In this article, I'll start off with the basics, provide you with several step-by-step procedures, and make sure that you have a good overall



Randy Williams

is a SharePoint MVP and senior consultant and trainer for Synergy Corporate Technologies, based in Singapore. He speaks about SharePoint topics at user groups and conferences.



Email



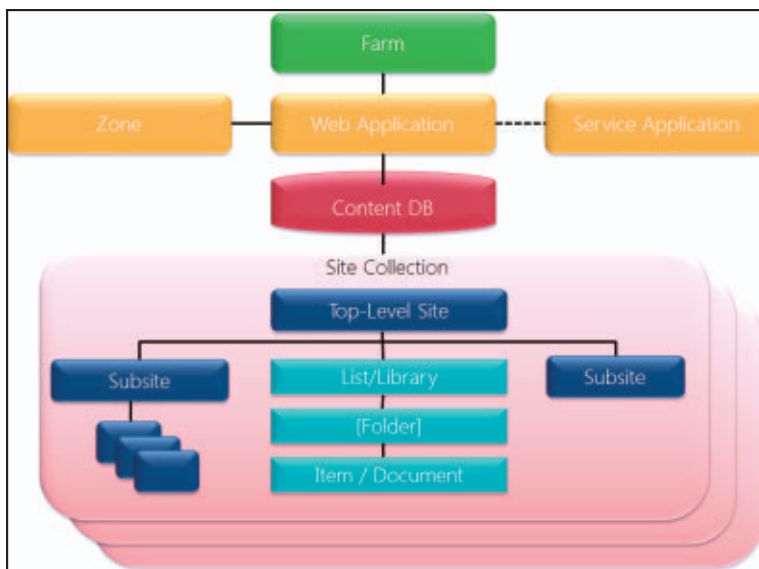
Blog

understanding of the SharePoint authorization model. The focus will be on the farm and work through web applications and site collections, down to individual sites, lists, libraries, and items. In Part 2, Kevin Laahs will focus on claims-based authentication; in Part 3, Todd Klindt will talk about hardening your SharePoint servers.

SharePoint's Logical Hierarchy

Before we get into setting permissions, let's review the basic terminology of a SharePoint environment. A SharePoint farm refers to all the servers (i.e., web front end, application servers, and database servers) that work together to provide a SharePoint service to users. Within the farm, a structured hierarchy organizes content, as Figure 1 shows. The purpose of this hierarchy is to organize and secure the vast amount of content within SharePoint.

Figure 1
SharePoint's logical
hierarchy



Securing the Farm

Despite the enormous responsibility that it implicates, securing the farm is quite a simple process. Permissions are mostly cut and dried: You're either a farm admin, or you aren't. A farm admin has full

control over the farm, including managing web applications, starting and stopping services, backing up or restoring the farm—effectively, any task that can be found in Central Administration, SharePoint’s web-based administration interface. A farm administrator can also grant anyone access to any and all content within SharePoint. Indeed, the farm admin role is a powerful one, and you should be cautious about who is granted this level of access.

Farm administrator access is granted by adding Active Directory (AD) users or security groups to the farm administrators group within Central Administration:

1. Within Central Administration, click Security and then select Manage The Farm Administrators Group.
2. Add or remove AD users or groups as needed.

Security note: To perform certain administrative tasks such as creating web applications, you also need to be a local Windows administrator on the web server (or servers) that run the Central Administration web application.

Granting PowerShell Access

Although Central Administration is commonly used for many day-to-day operations, Windows PowerShell is an incredibly powerful way to administer and automate SharePoint functions from a command-line interface. (To learn more about using PowerShell with SharePoint, see [“Using Windows PowerShell to Manage SharePoint 2010”](#) and [“SharePoint Administrators Can Learn to Love PowerShell.”](#))

To grant permissions for others to run PowerShell commands, use the Add-SPShellAdmin cmdlet. This command grants necessary permissions on SharePoint servers, the configuration database, and (optionally) a SharePoint content database. If you run this command without specifying a database, then the user is granted access to the farm only, not to any content databases. To grant a user the ability to run all cmdlets within a content database (e.g., to use Get-SPWeb to

return a site), use the -database switch to grant the user access to the content database.

Security note: To use the Add-SPShellAdmin cmdlet to grant PowerShell access, you must be a farm administrator, a local administrator on the server on which the command is run, and a security admin on the database server (or servers) to which you are granting permission.

Securing a Web Application

A web application is a user's entry point into SharePoint. Specifically, a web application consists of one or more Microsoft Internet Information Server (IIS) websites that control how users are authenticated (or not authenticated, if you allow anonymous access). Organizations commonly have multiple web applications to isolate certain types of content or sets of users, such as a Record Center or extranet, respectively. When granting permissions to a web application, the sole purpose is to control who has access to the content within the site collections that are associated with the web application. In other words, you cannot grant permissions to administer a web application; only farm admins can do that, and they have administrative control over all web applications.

To adjust permissions to content within a web application, change the user policy for the web application. You can make this change in Central Administration:

1. Within Central Administration, click Application Management.
2. Select Manage Web Applications.
3. Highlight the web application and click the User Policy button in the Ribbon.
4. Add or deny permissions, as needed.

Security Note: Because permissions can also be set for individual site collections (which I cover in the next section), permissions aren't usually granted this way. This method is very powerful and is used for special purposes. For example, you might decide to use

it to grant read-only access to auditors, to ensure that they have access to all site collections in the web application. Furthermore, this web application policy is the only method for denying someone access to SharePoint. This setting takes precedence over any other permission that is granted. Also, none of the permissions that you grant here are visible at a lower level. For example, a site collection administrator can't see who has been granted permissions through the web application policy. Again, this method of setting permissions is powerful but potentially confusing and dangerous—so use it with caution!

Allowing Anonymous Access

Anonymous access allows non-authenticated (i.e., guest) users to access a SharePoint web application. Anonymous access is commonly used for Internet-facing sites but can also be used for intranets. Anonymous access can be enabled whether you use claims or classic authentication. Here's how you enable anonymous access:

1. Within Central Administration, click Application Management.
2. Select Manage Web Applications.
3. Highlight the web application and click the Authentication Providers button in the Ribbon.
4. Select the zone. (A zone allows you to expose a web application with a unique URL. You can have as many as five zones per web application. All zones have access to the content databases and site collections that are associated with that web application.)
5. Select the Enable Anonymous Access check box and click Save.

Security note: By default, anonymous users have no access to site collections until that permission is granted. Therefore, after anonymous access has been enabled, you must set which permissions anonymous users have. (I cover this topic in the Grant Anonymous Access section of this article.)

Securing a Site Collection

Although you can grant permissions to multiple site collections through the web application policy, save this method for special circumstances. Users are authenticated at the web application level and authorization is typically performed at the site collection level. In other words, for regular, day-to-day permission management, you should assign permissions via site collections.

By default, permissions are inherited or cascade down to all web-sites, lists, libraries, folders, and items in the site collection hierarchy. Thus, permissions that are granted to the top-level site also apply to a document buried deep within a site collection. As you'll learn, this inheritance can be stopped.

The highest level of access within a site collection belongs to the site collection administrator. This user has implicit access to all SharePoint content (e.g., sites, lists, libraries, items) within the site collection, whether or not permission is specifically granted. Through Central Administration, a farm admin can specify up to two site collection administrators for each site collection. Within each site collection, however, a site collection administrator can add additional admins as needed.

You grant regular users (i.e., non-administrators) access through the site collection. In fact, each site collection is administered separately, making these collections very convenient administrative units that can be delegated out. Because administration occurs from within the site collection (and not Central Administration), administrative tasks are often assigned to trained business users, empowering them to be somewhat self-supporting. Of course, depending on your site collection architecture, this administrative boundary can become a challenge. For example, if you need to grant access to multiple site collections, that access must be granted one collection at a time for each user or group of users.

Security Note: If a user tries to access a site to which they don't have permissions, of course they get an access denied error. But you

can have the user request access from this error page. To learn more, see [“How to Manage SharePoint Site Access Requests.”](#)

Understanding SharePoint Groups

Like many other systems, SharePoint uses the notion of groups to simplify the assignment of permissions to multiple users. Note that SharePoint groups are specific to a site collection. Thus, you cannot use one SharePoint group to span multiple site collections. Again, this can be a blessing (when you want administrative separation) or a curse (when you don't). Also, SharePoint groups can't be nested, meaning that you can't place one SharePoint group within another. Some of these limitations don't apply to AD security groups, however. For example, you can assign one AD group permissions to multiple site collections. To better understand your options, here are the different ways permissions within a site collection can be granted:

- Add one or more users to a SharePoint group and grant that group permissions.
- Add one or more users to an AD security group and grant that group permissions.
- Add one or more users to an AD security group, add the AD group to a SharePoint group (this type of nesting is allowed), and grant the SharePoint group permissions.
- Grant permissions directly to one or more users, without using groups.

You assign permissions to users or groups by granting permission levels. Each site collection has its own built-in set of permission levels; you might be familiar with some of them (e.g., Full Control, Design, Contribute, Read). You can also create custom permission levels for more granular control over the granted access. Here are the steps to grant permissions within a site collection:

1. Within the site collection, go to the Site Actions menu and select Site Permissions.

2. On the Ribbon, click the Grant Permissions button.
3. In the dialog box, select individual users, SharePoint groups, or AD security groups.
4. Specify whether you want to add these users to a SharePoint group or to grant permissions directly. (Note: If you selected one or more SharePoint groups in step 3, you can only grant permissions directly.)
5. Optionally, send these users a welcome email.

When creating a new SharePoint group, you can also assign permissions to the group. Doing so makes permission assignment easy because you need only add users into the group. You can create new SharePoint groups from the same Site Permission screen on which you grant permissions.

Security note: When creating a new site collection, three default SharePoint groups are created automatically: owners (who have Full Control access), members (who have Contribute access), and visitors (who have Read-Only access).

Permission Inheritance

By default, permissions that are set for the top-level site within a site collection apply to all content in the site collection hierarchy (as Figure 1 shows). This concept is called permission inheritance and it helps to simplify permission management.

You can no doubt think of cases in which you need to establish unique permissions for a certain level. For example, you might want to protect a subsite and ensure that only managers have access. Within SharePoint, you can stop this permission inheritance on four different objects within the site collection: sites, lists/libraries, folders, and items. When you stop inheriting permissions, you can establish a unique ACL for that object, which then cascades to lower levels. This gives you a great deal of flexibility to mold a specific set of permissions around your content. However, be careful: The more you stop

inheritance, the more complex permission management becomes over time. Another troublesome side effect of stopping inheritance too often is that it can drastically affect SharePoint's performance. As a general rule, stopping inheritance should be the exception, and you are encouraged to structure your sites, libraries, and folders accordingly. Figure 2 depicts an example of permission inheritance.

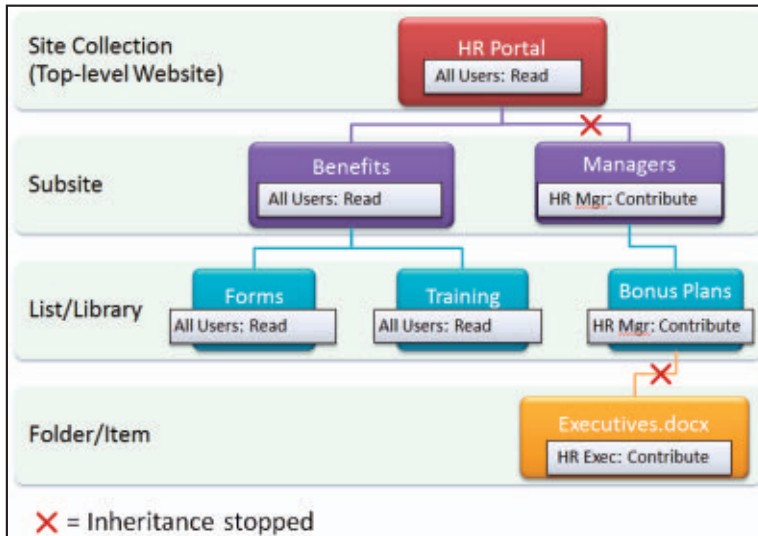


Figure 2

Permission inheritance
in action

The basic steps to stop inheritance and assign permissions are the same, regardless of the object type. For a site, use the Site Permissions link in the Site Actions menu. For a list or library, use the list's or library's settings page. For a folder or item, use the object's context menu, sometimes called the edit control block. If inheritance is stopped on an object, you can revert back to inherit, but any custom permissions that you applied will be lost.

Security note: Despite your best efforts, permission management can still spiral out of control. Another way to alleviate this is to have a good governance plan that gives administrative end users clear guidance about how to use groups and when or if to stop inheritance. Another option is to use third-party software such as AvePoint's

DocAve Administrator, which can help manage permissions across the whole farm in one operation. Tools like this also help enforce a security policy as defined in a governance plan.

Granting Anonymous Access

If the web application is configured to allow anonymous access, you can adjust which permissions an anonymous user receives:

1. Access the permissions page for the object, as described in the previous section.
2. If necessary, click the Stop Inheriting Permissions button in the Ribbon. (You can't adjust anonymous access on an object that inherits permissions from its parent.)
3. In the Ribbon, click the Anonymous Access button.
4. Choose the permissions to grant to anonymous users. The available permissions will vary depending on the type of object.
5. Click OK to save.

Wrapping Up

Hopefully, this security overview wasn't too painful and served as a useful refresher for those who have a SharePoint background. By now, you should understand SharePoint's authorization capabilities and some of its limitations. You've learned how to define farm and PowerShell administrators, adjust a web application's user policy, and configure each of the four permissionable objects (i.e., site, list/library, folder, and item) within a site collection. I've also given you some quick steps to enable anonymous access and set anonymous permissions. In the next article of this series, Kevin Laahs will focus on authentication, specifically the claims-based authentication support in SharePoint 2010. Until then, be safe and secure. ■

InstantDoc ID 143176

Product News for IT Pros

StillSecure Provides Easy Mobile Device Management

StillSecure announced upgrades to Safe Access, a NAC solution that lets you quickly and safely enable the use of personal mobile devices and tablets onto your corporate network with its bring-your-own-device (BYOD) identification and enforcement capabilities. With Safe Access 6.0, enterprises will also have the ability to identify not just mobile devices but also game systems, printers, and VoIP phones, and selectively allow or deny access to those endpoints. Compliant, optimized, and specialized for deployment in defense environments, Safe Access 6.0 will also provide information assurance vulnerability (IAV) reporting capabilities for its federal and Department of Defense (DoD) customers, including three of the four major military branches. By leveraging StillSecure's exclusive military-grade IAV rule set, customers can now work directly with IAV findings that are the lifeblood of DoD information assurance activity. For more information, visit the [StillSecure website](#).



SolarWinds Provides Comprehensive IP Infrastructure Management

SolarWinds released IP Address Manager (IPAM) 3.0, which gives you detailed visibility and management of your IP infrastructure, helping you better manage your network and minimize downtime. With the addition of centralized management of Microsoft DHCP servers and monitoring of Microsoft DNS and Cisco DHCP services, the company now offers a cost-effective way to manage the IP infrastructure. The number of IP addresses that you need to manage is quickly growing as a result of the increase in mobile devices, virtualized services, IP telephony, and other networkable endpoints. SolarWinds IPAM



is easy to deploy and use, addressing this continued growth of IP addresses for organizations of all sizes. Learn more today at the [SolarWinds website](#).



Veeam Offers Free Backup for VMware and Hyper-V

Veeam Software introduced the successor to its widely used free Fast-SCP tool: Veeam Backup Free Edition. The new free product, released in conjunction with Veeam Backup & Replication 6.1, offers ad-hoc backups of VMware and Hyper-V virtual machines (VMs). The features of Veeam Backup Free Edition address tasks that VM administrators deal with every day. VeeamZIP provides ad-hoc backup of a running VM for operational, archival, or portability purposes; Instant File-Level Recovery lets you restore individual guest files directly from an image-level backup; File Manager lets you easily manage VM and host files; and Quick Migration for VMware lets you migrate a running VM to any host or data store—even if you don't use clusters or shared storage. Learn more by visiting the [Veeam Software website](#).



F5 Enhances Suite of DNS Services with Added Scalability and Security

F5 Networks announced dynamic services to help organizations intelligently respond to DNS queries and application requests, consolidate infrastructure, and provide end-to-end protection for their essential DNS systems. With F5, customers can reduce internal and in-network response latency by up to 80 percent, seamlessly scale their systems up to 10x as demand grows, and provide advanced DNS security for physical, virtual, and cloud environments. F5 BIG-IP solutions are ideally positioned to intelligently manage and distribute DNS and application requests based on business policies, access demands, and network conditions—yielding an optimized and reliable experience for users as well as IT personnel. With this announcement, F5 is introducing new capabilities—including DNS caching and resolving, DNS Security Extensions (DNSSEC) validation, monitoring, and other

performance enhancements—building on previously announced solutions for Communications Service Providers and DNSSEC signing. For more information, go to the [F5 Networks website](#).

Symplified Introduces IDaaS Platform

Symplified announced Symplified Structure, which enables telcos, cloud service providers, cloud application brokers, and cloud application hosting companies to seamlessly integrate identity and access management (IAM) across their suite of cloud services. Symplified Structure is an instantiation of Symplified's core platform that unifies identity management, access control, auditing, and user provisioning across existing public/private cloud services and Software as a Service (SaaS) marketplaces. This fully embedded offering also allows service providers to deliver the convenience of secure single sign-on (SSO) from any device to support the skyrocketing use of personal mobile devices within enterprises. For more information, visit the [Symplified website](#).



ATEN LCD Supports DVI

ATEN announced a new single-rail LCD console with support for DVI that complements the company's existing full line of LCD console products. The new CL6700N DVI LCD console features an integrated 19" high-resolution LCD monitor, a full keyboard, and a touchpad housed in a 1U rack-mountable form factor. The CL6700N, which also supports audio, serves as the front-end sliding console for compatible DVI-based KVM switches. Companies that own compatible DVI-based KVM switches can take full advantage of the space-saving and efficiency benefits of the sliding console module without the need to purchase an additional KVM switch. The CL6700N LCD console offers connections for USB, VGA, or DVI devices at the console and server side and displays video resolutions of up to 1280 × 1024 at 75Hz. Utilizing DVI-I, the console supports both analog and digital video input. For more information, check out the [ATEN website](#).





Network Instruments Expands Its Insight

Network Instruments announced the latest release of its resource and infrastructure monitoring solution Observer Infrastructure (OI). OI offers significantly expanded visibility and monitoring of applications, as well as complex virtual and cloud environments. Through the integration of advanced flow technologies such as Flexible NetFlow, IPFIX, Citrix AppFlow, and Cisco Medianet Performance Monitor, OI queries systems, applications, and critical resources to provide advanced insight into the services and conversations traversing the infrastructure. IT teams can now utilize NetFlow and IPFIX within OI to obtain detailed traffic metrics valuable for troubleshooting, assessing bandwidth usage, and optimizing infrastructure performance. In addition, new product support for AppFlow allows OI to provide enhanced analysis of applications traversing Citrix virtual environments. Capturing AppFlow records generated by Citrix Netscaler application delivery controllers, OI tracks application status codes, average request response times, and average turn response times. All of these new metrics are then correlated alongside existing application and infrastructure data, providing a complete view of performance. For more information, visit the [Network Instruments website](#). ■

PAUL'S PICKS ▼

www.winsupersite.com



SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

Xbox Music (Preview)

PROS: Works with Windows 8, Windows Phone, and Xbox 360; online service gets key enhancements; Zune Music Pass will continue

CONS: Unclear if Microsoft plans to charge for this service (or make it part of Xbox LIVE Gold) and continue compatibility with legacy Zune devices

RATING: ★★★★★☆

RECOMMENDATION: Zune Music has always offered several key advantages over competing online music services from Apple, Amazon, and others, not the least of which is the excellent and affordable Zune Music Pass subscription service, which provides access to about 20 million songs. Microsoft is relaunching it as Xbox Music. Plans include expanding the library to 30 million titles and adding new native apps for Windows 8, Windows Phone 8, and Xbox 360, and a new cloud-based playlist. But Microsoft hasn't

come clean on what it will do with Xbox Music, and a poorly worded description of the service suggests that the company will charge for it by making users subscribe to Xbox LIVE Gold for \$60 per year. And it's not clear if legacy devices such as Zune will be compatible with the revamped service. But it looks solid for music lovers betting big on Windows 8/Windows RT, Windows Phone 8, and the Xbox 360.

CONTACT: [Microsoft](#)



Full Review

Xbox 360 SmartGlass (Preview)

PROS: Lets Apple iPhones and iPads, Android devices, Windows Phones, and Windows 8/Windows RT tablets interact with Xbox 360 in unique new ways

CONS: Capabilities are unclear; won't be universal across content types

RATING: ★★★★★☆

RECOMMENDATION: Xbox SmartGlass was one of the most widely touted and least-understood Microsoft announcements at the E3 trade show this year. It's basically a mobile app and associated platform and the 2.0 version of an existing app called Xbox Companion. It allows compatible mobile devices to act as a remote control for the Xbox 360. You can also use it to discover games, TV shows or movies, and music to play on the console. What's new in SmartGlass, aside from the cross-platform support, is that it's a mini-platform of sorts, letting you use the device as a semi-connected second screen. Whether SmartGlass lives up to its lofty prerelease hype remains to be seen, but even if it's just a warmed-over Xbox Companion, that's not too shabby either.

CONTACT: [Microsoft](#)



Full Review



Best of TechEd 2012 Winners

Recognizing the best products from Microsoft and its partners



Jason Bovberg

is a senior editor for *Windows IT Pro* and *SQL Server Pro*, covering products as well as the systems management, hardware, and storage/backup topic areas. He has more than 20 years of experience as a writer and editor in magazine, book, and special-interest publishing.

Email



Twitter



Website



At Microsoft TechEd 2012 in Orlando, Florida, our editors honored this year's Best of TechEd Award winners. These awards recognize Microsoft partners that offer innovative products and services in the marketplace. The judging panel—consisting of *Windows IT Pro*, *SQL Server Pro*, *Dev Pro*, and *SharePoint Pro* editors Michael Otey, Sean Deuby, Megan Keller, Jeff James, and Jason Bovberg—narrowed an impressive field of 260 submissions down to 46 finalists in 15 categories. On the show floor, the team interviewed the finalists and evaluated the products to determine a final list of winners. As always, the three criteria for the judging process were strategic importance, competitive advantage, and value to customers. Show attendees also cast their votes to determine the winner of the prestigious Attendees' Pick award.

It was a fun week in Florida, capped off with a wonderful awards party at the Universal Studios Islands of Adventure theme park. We would like to congratulate our 2012 winners!

Backup and Recovery

Symantec's NetBackup 5220 Appliance



Symantec's NetBackup 5220 appliance stands out among a growing field of dedicated devices that are purpose-built as all-in-one backup solutions for physical, virtual, and cloud-based data protection.

Engineered for midsized-to-large enterprises, remote offices, and data centers, the Intel-designed NetBackup 5220 appliance truly shines in its virtualization capabilities, “green-friendly” non-CPU-intensive deduplication capabilities, and no-nonsense licensing.

Business Intelligence

Attunity Replicate

[Attunity Replicate](#) simplifies the tedious process of loading data from a source database into a target SQL Server database. Its Click-2-Replicate drag-and-drop functionality lets DBAs perform the loading task without having to write any code, and its dashboards let you monitor the entire process and easily see what data has been successfully replicated to the target database and what data hasn't. Attunity Replicate stands out for its ease of use, its zero footprint, and the fact that it enables data to be replicated quickly so that business users can make informed decisions based on that data.

Cloud Computing

Riverbed Technology's Steelhead Cloud Accelerator

[Riverbed Technology's Steelhead Cloud Accelerator](#) is a one-of-a-kind Software as a Service (SaaS) application accelerator for Office 365, Google Apps, and Salesforce.com that provides dramatic performance gains compared with non-optimized WAN connections. By making a remote cloud resource appear to be a fast local network resource, Steelhead Cloud Accelerator eliminates the performance bottleneck at the very heart of public and hybrid cloud computing.



Hardware and Storage

X-IO Technologies' Hyper ISE

[X-IO Technologies' Hyper ISE](#) is an extremely high-performance storage system that fuses traditional hard disks and SSDs into a single pool of capacity. Its new X-IO Orchestrator software is the software that elevates this solution into the stratosphere, providing real-time

provisioning of workloads to the correct disk resources. The performance numbers of the X-IO Hyper ISE continue to skyrocket, blowing away the competition in all kinds of real-world data-intensive applications and environments. This year, the X-IO Hyper ISE evolves from last year's Breakthrough Product to this year's Best Hardware and Storage Product.



Messaging and Unified Communications

Exclaimer Mail Disclaimers

[Exclaimer Mail Disclaimers](#) for Exchange 2007 and 2010 provides powerful tools to centrally manage and control email signatures and email disclaimers. Given the ubiquity of email signatures, Mail Disclaimers gives admins a powerful assortment of templates and tools to manage them from one central location. Because of the regulatory and compliance needs some organizations operate under, email disclaimer automation can be an enormous timesaver and headache-remover.

Microsoft

System Center 2012

[Microsoft System Center 2012](#)'s suite of virtualization and private cloud management tools is a compelling combination of both capability and value. This suite helps IT pros move their companies toward cloud computing and advance their own skills—at a time they most need it.



Mobile and Wireless

Nokia Lumia 900

The [Nokia Lumia 900](#) is unarguably the best Windows Phone ever made, with a bright 4.3" AMOLED display, dual cameras with Carl Zeiss optics, video calling, 4G LTE, and a 1.4GHz processor. More important, the Nokia Lumia 900 is the first Windows Phone that can legitimately compete head-to-head with the latest Apple and Android



smartphones. A decade from now, the Nokia Lumia 900 will be remembered as one of the most significant developments in the ongoing evolution of the Windows Phone platform.

Networking

Citrix Systems' NetScaler 10 with Citrix TriScale Technology

[Citrix Systems' NetScaler 10 with Citrix TriScale Technology](#) is available as both a physical network appliance and a virtual appliance. It sits in front of your web server, database server, or enterprise apps to provide load balancing, data compression, enhanced visibility, performance acceleration, high availability, and increased security. NetScaler provides particular benefits to SQL Server, optimizing its performance and capacity. The addition of TriScale to NetScaler brings a new level of scalability to the solution.



Relational Database

Quest Software's Spotlight on SQL Server Enterprise

Quest Software's SQL Server diagnostics tool, [Spotlight on SQL Server Enterprise](#), provides DBAs with valuable information about how their entire SQL Server environment is functioning by monitoring and diagnosing problems on not only their SQL Server systems but also their related Windows Server and VMware systems, as well as the SQL Server Analysis Services (SSAS), replication, and SQL Azure connection types. This product features a bevy of stand-out capabilities, including mobile device support, a Heat Map to prioritize issues, and diagnostics information that offers up the most common scenarios that might cause a specific problem.

Security

TITUS Message Classification

[TITUS Message Classification](#) is a user-driven email security tool that can be an invaluable asset for organizations operating under strict auditing,

compliance, and security requirements. TITUS Message Classification adds the ability for users to easily categorize email messages into required categories, and it helps ensure that sensitive email is delivered only to intended recipients. It integrates seamlessly with Microsoft Outlook, helps prevent accidental disclosure of sensitive messages, and can increase the effectiveness of an existing security, DLP, or archiving solution.

SharePoint Administration

Colligo Briefcase Pro for iPad

[Colligo Briefcase Pro for iPad](#) lets iPad users securely access their corporate SharePoint sites in an easy-to-use, Dropbox-like manner. It solves an immediate, pressing need to provide easy, yet secure, mobile access to IT-sanctioned resources instead of unmanaged external storage services such as Dropbox.

SharePoint Development

K2 blackpearl

[K2 blackpearl](#) lets SharePoint professionals automate business processes and create SharePoint workflows using their visual designer with no need for coding. Workflows created with the graphical K2 Designer or K2 Studio tools can automate almost all SharePoint tasks, including creating sites, managing documents, and prompting for process approvals. K2 blackpearl can be integrated with a number of external data sources, including SAP, CRM systems, and databases such as SQL Server.

Software Development

Telerik Dev Tools for .NET Ultimate Collection

The [Telerik Ultimate collection](#) is a full-featured suite of development tools and controls for ASP.NET, AJAX, Silverlight, MVC, WinForms, WPF, HTML, and JavaScript, as well as tools for reporting. The Telerik Ultimate collection also includes the company's JustCode code analysis and refactoring product and its free JustDecompile product, which can decompile .NET executions into source code.



Systems Management and Operations

SolarWinds' Server & Application Monitor

SolarWinds' [Server & Application Monitor](#) is a comprehensive systems and application management and monitoring solution that's capable of monitoring Windows and Linux and more than 2,700 different applications. Server & Application Monitor contains expert guidance about what to monitor, optimal operating thresholds, and best practices.



Virtualization

HP VirtualSystem for Microsoft

The [HP VirtualSystem](#) is an appliance that removes the complexity of implementing high-performance and scalable virtualization in the enterprise. This is a preconfigured appliance that has been expressly designed by HP and Microsoft to support enterprise-level virtualization. Proving that point, this year's 11,000 TechEd labs were all run on virtual machines (VMs) supported by the HP VirtualSystem VS3.

Breakthrough Technology

Cisco UCS Server and UCS Manager

[Cisco UCS Manager](#) provides complete programmability for all low-level hardware, BIOS, and configuration settings for Cisco UCS Servers, enabling them to be quickly deployed, cloned, and managed—even remotely. Cisco UCS Servers can also be fully managed using PowerShell or System Center Orchestrator.

Attendees' Pick

X-IO Technologies' Hyper ISE

For the second year in a row, the [Hyper ISE](#) takes the Attendees' Pick award. This is a truly exciting product—providing unheard-of performance and intelligent workload-provisioning management, not to mention a striking design. This powerhouse has it all: brains, beauty, and brawn. ■

InstantDoc ID 143449

ControlPoint 4.5



**Russell
Smith**

is an independent IT consultant specializing in systems management and security, and author of *Least Privilege Security for Windows 7, Vista and XP* (Packt).

Email



Twitter



Axceler's ControlPoint 4.5 is a management solution that provides a means to manage SharePoint security, configure sites, plan for storage capacity, and generate reports on SharePoint content. Tightly integrated with SharePoint, ControlPoint is a SharePoint web application that can offload the statistics and data it collects to a SQL Server instance running on a server that's not hosting SharePoint.

Installing ControlPoint

ControlPoint supports all editions of SharePoint 2010 and SharePoint 2007. As a SharePoint web application, ControlPoint must be installed on a server running all SharePoint roles or a SharePoint web front-end server. Although this requirement might seem like a potential disadvantage because of the increased complexity and load on SharePoint, it gives users access to ControlPoint's advanced management features through SharePoint, which provides a familiar and easy-to-use interface.

ControlPoint uses two databases: a content database that stores SharePoint content and a service database that stores administrative information and data collected from SharePoint. I installed ControlPoint on a server running SharePoint Foundation 2010 and used the instance of SQL Server 2008 R2 that was installed automatically with SharePoint Foundation 2010 to host the content database. I installed the service database on a separate server running SQL Server 2008 R2.

ControlPoint requires SQL Server 2008 R2 or SQL Server 2008 for its content and service databases when used with SharePoint 2010. If you want to use ControlPoint with SharePoint 2007, the database requirements differ slightly. SQL Server 2000 through SQL Server 2008 R2 can be used to host the content database, but only SQL Server 2005 through SQL Server 2008 R2 can be used to host the ControlPoint service database. Other system requirements are Microsoft IIS and Microsoft .NET Framework 3.5 SP1.

To meet the necessary permission requirements, I used a domain account that had local administrator permissions on the SharePoint server, sysadmin permissions on my SQL Server 2008 R2 server, and SharePoint farm administrator permissions to both launch the installer program and specify the ControlPoint service account during the installation process. Installation went smoothly, with a prerequisite check identifying only one issue. The SharePoint Administration service wasn't running on my SharePoint server, which was easily rectified.

After the web application is installed on the SharePoint server, there's an activation process that involves entering a license key and letting ControlPoint activate the installation over the Internet. The license information must then be recorded by clicking a separate button, and the installation is updated and verified as activated.

Navigating ControlPoint

Despite ControlPoint running in a web browser, the rich UI lets you right-click items in the SharePoint hierarchy and select options from a context menu. This feature isn't something you always see in web-based applications, but it makes ControlPoint feel much more like a standard desktop program and makes its UI easier to use.

When you start ControlPoint for the first time, there's a message at the top of the UI informing you that a full discovery has never been performed. Unfortunately, it's not an active hyperlink so you need to know where to look to run a discovery operation. It's not hard to find, however; it's under the Manage ControlPoint tab in the Outlook-style panel on the left. Discovery is a scheduled task, but there's the option to run the task immediately. You can monitor the task's progress using the timer at the top of the window.

Using ControlPoint to Copy Sites and Site Collections

ControlPoint doesn't reinvent the wheel. Where SharePoint provides adequate administrative capabilities, ControlPoint hands you over to SharePoint to complete the task. One area in which ControlPoint

**ControlPoint
doesn't
reinvent
the wheel.**

provides welcomed additional functionality is when you're copying sites and site collections. Unlike native SharePoint administration, ControlPoint lets you manage sites, web applications, and site collections across an entire farm in a single operation.

The first task I attempted was to copy the default Team Site. The copy operation is in the Content context menu. When selected at site level, you're given the option to copy the whole site and all its content, including subsites. There's also an option to exclude child sites or copy just the site configuration without the content. In a separate window, I specified the URL for my new site and the path to the file in which the source and destination farms could write temporary files. You also have the option to specify how security should be applied to the new site (i.e., inherited from the parent or taken from the source site). Before the copy operation is run, a comforting Verify Request Action window is displayed with a summary of the selected operations.

After the copy operation has completed, you're automatically taken to the results window. Interestingly, there's an option to save the instructions to an .xml file, which you can later import when you need to repeat the same operation in the future. Deleting a site is easy, and you have the option to export the site contents before deletion.

When copying or moving a site or site collection in a multi-farm environment, the source selection is limited to the home farm but the destination can be in any farm. Also, copy and move operations can't be performed between SharePoint 2010 and SharePoint 2007 farms. Axceler has a separate product ([Davinci Migrator](#)) for SharePoint migrations.

ControlPoint's search function is limited to locating sites in your SharePoint environment, but the search can be based on advanced criteria. If you have more than one SharePoint farm, ControlPoint lets you run a search across multiple farms by selecting them in the left pane, right-clicking, and selecting Advanced Search in the context menu.

Using ControlPoint Policies and Groups

Using ControlPoint policies, you can control the SharePoint environment. For example, you can create policies that restrict lists, set quotas to limit storage on disk, and notify users when they exceed those quotas. ControlPoint policies are limited to the home farm in a multi-farm environment.

When creating a policy to restrict lists, you can select all lists or a specific type of list. Similarly, you can restrict all users or specific users by SharePoint group, Active Directory (AD) group, AD user, or permission level (e.g., Full Control, Contribute, Design).

There's comprehensive management for permissions, including the ability to create group associations. With the group associations, you can propagate membership and permissions from a model group to any chosen dependent groups.

You can configure object properties (e.g., settings for creating minor and major versions of documents) in bulk across an entire farm. With the standard SharePoint administration interface, such configurations would have to be completed manually for each object. Further, in ControlPoint, you can enforce settings and make them the default for new document libraries. There's a comprehensive list of auditing options for site collections, such as editing users and permissions, and checking items in or out. You can even create alerts to notify you when objects are modified.

Access to menu items in the ControlPoint interface can be controlled by ControlPoint groups. Access can even be given to ControlPoint features such as links on SharePoint pages, enabling standard users to get access to ControlPoint so they can administer their own sites without having to use the ControlPoint interface.

Using ControlPoint to Analyze Data

Data collected in ControlPoint's service database can be analyzed, thus avoiding the need to put a high load on your SharePoint environment to collect the necessary information. You can see the most

ControlPoint 4.5

PROS: Tight integration with SharePoint; well thought-out administration tool

CONS: Some functions restricted to the initial SharePoint farm added to ControlPoint; no rollback functionality

RATING: ★★★★★

PRICE: Based on the size of the SharePoint environment (e.g., a license for one farm with multiple SharePoint web front ends and 5,000 users is \$18,000)

RECOMMENDATION: If you need to manage a large SharePoint estate, standardize settings, and get a grip on security, ControlPoint is a good solution. It provides the necessary tools for administrators and SharePoint users alike.

CONTACT: [Axceler](#) •
866-499-7092 or
781-995-0063

or least activity for site collections, sites, lists, and list items. You can also check the activity level for documents and users. Analyses can be run over a given period for the purpose of identifying trends, although there are some restrictions if you're running a Windows SharePoint Services (WSS) based environment.

To run an analysis, you select the object (or objects) in ControlPoint, right-click, select Activity, and choose Activity By User or Activity By Document. Using the Interactive Analysis feature, you can work with a data set from a search in tabular format and create different types of charts. ControlPoint includes custom lists for farm, site collection, and web application statistics. You can use these custom lists to create SharePoint web parts that form the foundation for SharePoint dashboards. You can make the dashboards visible to users, administrators, or both.

On the Ball

ControlPoint offers a lot of advanced features, such as the ability to copy sites and site collections, and use existing sites as templates. These features will be especially valuable for large organizations that need to adhere to industry regulations and ensure that security is configured to a known standard across the SharePoint estate.

Solutions that offer many advanced features can be complex to work with, but that's not the case with ControlPoint. Everything is very logically laid out, and it's easy to find exactly what you want without having to scratch your head and dig deep into the UI. The documentation is also good.

The one major feature that I'd like to see in the product is an undo function. Although every action is logged, an undo function would give administrators confidence that, should the worse happen, there's a quick and easy way to roll back changes. ■

InstantDoc ID 143318

Avance

Floods, fires, earthquakes, power outages, and software and hardware failures are reminders of why disaster readiness and recovery are so important. Maintaining business continuity in the face of this adversity could mean the difference between weathering the storm and going out with the lights.

Enterprise IT groups know and handle this challenge well, but it can be quite difficult for smaller organizations to meet 99.0 percent uptime requirements, let alone 99.999 percent. Cost and complexity barriers keep many businesses from trying high-availability solutions at all, forcing IT staff to use manual, administrator-intensive detection, remediation, and recovery processes.

Many forms of high-availability solutions exist today, ranging from software-based solutions to mission-critical solutions that offer hardware-level redundancy and failover. The trick is to pick the right one for your organization, thereby achieving the desired availability without breaking your IT budget. As with network security, the more you can afford the better off you'll be, but there is a tipping point at which you're throwing good money after bad. In other words, your particular business might not require extreme measures. I recently took a look at Stratus Technologies' Avance high-availability software, one of the midrange solutions that can deliver availability at near-enterprise levels, but without the million-dollar outlay.

Overview

CIOs often call on systems administrators to reduce costs but still boost IT reliability. Administrators in small-to-midsized businesses (SMBs) tend to feel this crunch more acutely, because delivering fault tolerance can more than double the cost of the existing infrastructure for backup servers, redundant networking, and so on. Although native technologies in Windows Server are capable of getting you part



Joel Sloss

is formerly the lab manager and technical editor for *Windows IT Pro*. He has been in the enterprise IT software industry for more than 17 years and has authored numerous articles and contributed to several IT books.



Email

of the way there, they fall short of the instantaneous failover that's needed for demanding workloads—and demanding CIOs.

Stratus aims to solve this conundrum through a hardware agnostic, yet not entirely hardware independent, software-based availability package for SMBs. Stratus has made its name in enterprise-class high-availability solutions for more than 30 years, keeping the lights on 24 × 7 for critical human services, such as 911 call centers, hospitals, utilities, and more.

Avance combines a software offering with proactive management (which can even be monitored by Stratus remotely) and hardware redundancy. An Avance high-availability cluster provides near-zero failover and recovery times, with near-zero client impact (including stateful applications) using real-time monitoring and data replication. If you're running a heterogeneous environment, you'll also appreciate Avance's support for Linux server platforms (e.g., Red Hat, CentOS) and applications. Avance uses CentOS 5.5 and Citrix Systems XenServer virtualization technologies to abstract hardware from software, providing a foundation for transparently migrating OS and application workloads between physical systems in the event of a failure.

You can use most of the off-the-shelf server, networking, and storage hardware as long as any two systems you cluster are similar enough that a hardware mismatch doesn't result in bad driver behaviors (and thus a crash). In addition, the same RAID configuration must be used on both machines. One benefit of this clustering approach is that you don't need to purchase a dedicated storage array for data because replication between servers occurs over the wire.

The downside is that you still need an equivalently configured second server as a hot standby. Note that you won't have an active-active performance cluster. For more information, see the sidebar "How Avance Works."

Setup

For expediency, I started with two white-box Intel servers, which were supplied by Stratus. Each server had a S5520UR motherboard,

How Avance Works

It's somewhat difficult to categorize the high-availability model used by Stratus Technologies' Avance software. It isn't really a sharing, active-active, or active-passive model but rather has elements of each one. It could be likened to an instantaneous hot standby, perhaps similar to a RAID 1 storage volume. For this reason, I'm going to call it a "mirrored cluster." Data is replicated in real time from the primary node to the backup node over a private link, just as an array controller replicates a mirrored volume's I/O to both drives in a logical disk pair. In Avance, however, only one node is fully operational and running the software at a time, whereas in a RAID 1 volume, both disks are used simultaneously.

In the Avance cluster, both machines are aware of the system state and data state at all times (down to the microsecond) because what happens on one node immediately happens on the other node in exactly the same way. Failover between nodes is instantaneous without missing a heartbeat (sort of like a parallel universe). Because the second node is running essentially in lock-step with the first, if the cluster's virtual IP switches over to the backup node, neither the application nor the client are aware anything happened. Not a single bit is lost in translation, and not a second is wasted.

In contrast, an active-passive cluster requires the secondary node to realize that the primary node has failed, take ownership of the shared resources (usually disks), recover application state, and reinitialize connections. Transactions can be lost, triggering further recovery processes within an application database. Plus, there's a necessary time lag for this all to occur.

As with other types of clustering, clients connect to a single virtual IP address that's managed by the Avance software, which manages physical links and IP addresses. The management console and all workloads run on only one physical machine at a time. However, instead of a situation in which an active connection on the failing node would be lost (requiring the client to reestablish the connection), memory state is maintained and the session (such as Remote Desktop Services) isn't affected. Workloads can flip back and forth between nodes as many times as you want, without impacting service availability. ■

InstantDoc ID 143591

dual quad-core Xeon X5560 processors, 24GB of memory, and 2TB of disk space. You can gain additional hardware resiliency if you select a chassis with hot-swappable components (e.g., CPU, RAM), RAID controllers, redundant power supplies, failover NICs, and so forth; doing so will reduce the likelihood of a single-server failure. This isn't required, however, since the solution's real-time monitoring includes more than 150 different metrics and predictive analytics that will trigger a live migration if a fault is either detected or about to take place.

Your dual-server configuration doesn't need to be any different from your standard build, with the exception of a dedicated gigabit Ethernet port on each machine for management and data replication, which is referred to as the "Sync" link. The servers can also be completely headless (after initial setup), because all maintenance operations are performed through a web-based console. However, Stratus recommends redundant Sync links to improve performance and fault tolerance.

Avance installation is straightforward and uses a self-imaged DVD. It automates setup for both servers through a single process, but you should reformat the machine if you're repurposing older hardware. (You can't change out the hardware on an existing OS platform build or migrate it from another machine unless it's identical hardware and already virtualized.) Adding the second machine to form a cluster is achieved by a fast software install driven from the primary node. When you join the second server to the cluster, an automated synchronization process images and configures it.

Avance's instant data replication between nodes means each server is always up to date. When a hardware failure, predicted failure, or planned shutdown occurs, the second machine simply picks up where the first one left off. This lets you carry out whatever maintenance is required on the first node without a service interruption. When you're done, you can manually flip the workload back to the first node or leave the workload on the second node, letting it migrate back to the first node only if a failure is detected on the second one.

Operations

Avance features a web management console (Apache Tomcat based on HTML5 and JavaScript), which runs on only one cluster node at a time. The great thing about web-based management is that it's usually simplified and available from any client (even a mobile device). The downside is the nagging question about security. Although you benefit from a thin-client (no install) experience, you're increasing

your potential attack surface by running—at the host level—a web application that is capable of full system control, even over SSL.

Each virtualized server workload can be locked down and protected with anti-malware solutions and the like, but host-level intrusions are bad news. XenServer isolates virtual machines (VMs) from each other, and Stratus has invested in hardening the host configuration.


For additional protection, you could deploy a full application-layer firewall and place your servers in a demilitarized zone (DMZ), which is a common topology. Alternatively, you could deploy a dedicated security VM through which all traffic gets routed. However, a bare-metal hypervisor with no native application operations and a separate management server would provide a better overall (albeit more costly) security posture. The Avance console has an inactivity time-out feature but lacks any token-based or multi-factor authentication capabilities.

Every task in Avance is possible through the web management console, saving you from having to sit in the wiring closet with the machines. It's also convenient if you employ a services management vendor to remotely maintain your IT infrastructure. Although native IP repudiation or filtering isn't supported for limiting administration through only certain IP addresses, access through a VPN and firewall will grant similar security.

Using an easy-to-understand layout, the UI gives you quick access to Avance's default dashboard, which provides alerts, configuration details, and drilldown pages for managing both physical and virtual cluster attributes. You also have quick access to pages in which you can manage physical machines, set up storage groups and volumes to dedicate resources to specific workloads, lay out virtual networks, manage users, build VMs, and more. Most operations are driven by easy-to-use wizards that automate the tasks.

In keeping with the fully virtualized nature of the solution, you can create virtual CD installation points accessible by specific VMs, as Figure 1 shows. They can be used as either direct copies of .iso

Figure 1
Creating virtual CD
installation points



software media or downloadable installs by both servers and virtualized desktops. Although it might not be advisable from a security point of view, you can make physical components such as USB storage available to individual workloads.

Failover

When a failure occurs, Avance provides active monitoring across a variety of different categories, enabling a full range of fault detection, whether physical or virtual. As with some out-of-band (OOB) management solutions, predictive filters can help identify when something bad is about to happen, instead of just waiting for a failure. With this fair warning, you can get ahead of the problem before a catastrophic event occurs that even Avance can't handle. Although if you're using the right combination of metrics, which are dependent on the specific hardware and OS, I'm not sure what this could be.

To test Avance, I did a number of disagreeable things to the servers. I removed network cables, unplugged the power cord, killed VMs,

**To test Avance,
I did a number
of disagreeable
things to the
servers.**

and so forth. I even went so far as to hard-crash both machines at the same time by yanking out all power cords, even to the redundant power supplies (causing them to emit a variety of plaintive beeps). Impressively, nothing bad ever seemed to happen. Killing one entire server produced a warning in the console, as Figure 2 shows, but neither the management application nor the workloads (such as the Remote Desktop Services session) seemed to notice.



Figure 2

Checking the warning in the default dashboard

The VMs seamlessly kept going. When I brought the failed primary server back online, it quietly rejoined the cluster, resynchronized its data, and took its place as the new secondary node. I had difficulty thinking of anything else I could break without physically damaging the hardware.

Given these capabilities, what could you use Avance for, beyond the obvious uptime enhancements? As I previously mentioned, there are other forms of fault tolerance and clustering available, some of which might be better suited to certain workloads or situations. Areas in which Avance would be a natural fit include:

Stratus Avance

PROS: Provides enterprise-class availability for SMBs; easy to configure and administer; fully automated; near-zero failover latency; versatile for heterogeneous software environments (supports Windows and Linux)

CONS: Doesn't support managing more than one cluster simultaneously, possibly making it difficult to manage multiple remote sites; operating environment is security-hardened but host-level web management application introduces potential attack surface

RATING: ★★★★★

PRICE: \$5,000 for dual-server Avance license, plus \$100 monthly maintenance (contact Stratus for more detailed pricing information)

RECOMMENDATION:

Avance is well-suited to SMBs or managed remote sites that require high availability. Avance supports most enterprise-class applications (e.g., Exchange Server, ERP, CRM software) and limited-scale database environments.

CONTACT: Stratus Technologies •
800-787-2887

- 99.99 percent application availability
- Remote-site redundancy
- Small or branch-office resiliency
- Small- to average-size workloads (e.g., Microsoft Exchange Server, Microsoft SharePoint, customer relationship management—CRM—software, limited-scale database environments)
- Private cloud

Areas in which a different approach (or perhaps the more advanced enterprise-class V Series offering from Stratus) would be best include:

- High-throughput transaction processing
- Data warehousing
- Real-time computing
- High-capacity distributed applications or enterprise-scale deployments (e.g., multi-server email or database environments)
- Public cloud

Note that there isn't a facility for managing multiple Avance deployments through a single console. Thus, building one large cluster of powerful machines would be better than using several smaller clusters in a demanding environment.

Avance Lives Up To Stratus' Reputation

Avance lives up to the reputation established by Stratus' more advanced availability solutions. Avance also provides capabilities you'd normally expect in much higher-priced packages. With its focus on failover and ease of use, smaller IT shops with limited resources or training will be able to up-level their service offerings and greatly enhance disaster readiness.

But perhaps a more important question might be, "Would I install this in my data center?" The answer is yes, I would. ■

InstantDoc ID 143501

Apple iPad for the IT Pro

Unless you've been living under a rock or disconnected in some other way from the rest of the world, you know that Apple's iPad has become a mainstay in the technology world. No matter where you go it seems there's someone toting an iPad, talking about new apps for their iPad, or heading off to buy an iPad. The extreme popularity of the device is highlighted by not only Apple's own sales numbers but also the media frenzy that occurs anytime even a hint of a new model hits the rumor circles.

Despite the device being capable of producing content, a large amount of time can be spent using an iPad to consume content. No matter whether you're browsing the web, watching movies, reading books, or playing games, iPad is an ideal content-consumption device. But what about using it for "real work"?

It's not too hard to find a plethora of articles and blog posts on both sides of the fence when searching the Internet. Some people say that there's no way they could replace their primary computing device, be it a desktop or laptop, with any kind of tablet—iPad or otherwise. "It's impossible to do precision work! The mouse is far more accurate," they cry. "I need a real keyboard. Can you imagine typing a report or working on a 30,000-row spreadsheet on an iPad?" people bemoan. Other people have no problem using an iPad as their primary device. "I only go back to my 'real' computer when necessary. The iPad is perfect for 95 percent of what I have to do," they explain.

I don't stand firmly in either camp. I believe in using whatever is best for the task I'm trying to accomplish. Right now I have a Lenovo ThinkPad running the Windows 8 Consumer Preview, an HP desktop running Windows 7 on a 23" monitor, Apple's Mac mini running OS X Lion on a 21" monitor, Amazon's Kindle Fire, and a third-generation iPad (see Figure 1). By trade, I'm a systems engineer. Can someone like me



Michael Dragone

is a contributing editor for *Windows IT Pro* and a senior network engineer. He holds MCDST, MCSE: Messaging, MCTS, and MCITP credentials and remembers when *Windows IT Pro* was called *Windows NT Magazine*.



Email

Figure 1
Third-generation iPad



successfully use an iPad for “real work?” And even if I can, would I want to? Let’s find out.

Putting It to the Test

The HP desktop is my primary device, and I have a fair number of applications installed on it for work. In addition to what you would typically find installed (e.g., Microsoft Office), I installed the Microsoft Lync client, Remote Desktop Connection Manager (RDCMan), PuTTY to take care of Telnet/Secure Shell (SSH)

jobs, the Remote Server Administration Tools for Windows 7, every web browser known to the world, and so on. I don’t believe I perform any task that’s so massive that it would be downright impossible on an iPad, nor are any of my applications in the category of “Whoa! What is that?” In fact, I installed many of the same tools on my iPad. There’s the Lync client, a Remote Desktop client (HLW Software Development’s iTap RDP), and an SSH client (Panic’s Prompt). I use the built-in VPN client to connect back to corporate resources. I also take advantage of multi-device Web 2.0 tools, including Evernote and Dropbox, to have access to just about anything I need from anywhere.

Perhaps proving that Apple’s marketing machine was right all along, procuring those iPad applications was a breeze. I simply tapped the App Store icon and searched for an application by entering its name. When the first-generation iPad was new to the market, I would typically search by entering what I was looking for (e.g., “SSH client”). Now it’s even easier to find iPad applications that meet your needs, as new apps are constantly being reviewed on the Internet. Try searching for “RDP client iPad” and you’ll see what I mean.

For the most part, configuring the applications was also a breeze. For example, I didn’t have to perform any “iPad trickery” to get iTap

The iPad is an ideal content-consumption device, but what about using it for “real work”?

RDP to connect to my Remote Desktop Session Host server, nor did my Cisco VPN require any configuration changes. I was pleasantly surprised to find that all of the applications that connected back to corporate resources did indeed “just work.”

With all that being said, you might think utilizing an iPad for “real work” is nothing but roses. However, after spending a few days attempting to use my iPad for my “real work,” I found one drawback: On the iPad, the only pointing device is your finger and so much of Windows administration is still GUI-based.

Connecting to corporate resources with the VPN client isn’t an issue. Accessing internal websites, even internal websites used for administrative tasks, is fine (as long as the site doesn’t require a plug-in such as Microsoft Silverlight or Adobe Flash). Getting to a router, Linux host, or UNIX host is no problem. Even using a tool such as LogMeIn to see what someone is doing on a machine and talk them through a problem isn’t a challenge.

However, using Remote Desktop to administer a Windows server is a challenge. If you’re thinking that you can simply set up an SSH connection to a Windows host and run command-line utilities from there, stop and think about all of the server applications you have installed on your Windows servers. Then think of the applications that just stink in terms of remote administration. I’m not referring to remote administration from a mobile device while you’re in the jungle; I’m referring to remote administration that stinks even from your desk 100 feet away from the server, such as:

- The fax server that’s administered from a GUI-based console that hasn’t changed since 2001
- The internally developed server application that runs only on Windows Server 2003 and works only by using the application’s console interactively
- The telephony application in which the only command-line access is to start and stop it, with no capability to look at and troubleshoot log files

For these types of server applications, you need Remote Desktop. Even with pinch-to-zoom, three-finger swipe, and all of the other niceties that are built in to applications like iTap RDP, it's still a challenge to use Remote Desktop for Windows server administration because of the GUI elements and the lack of a mouse. If I only had to administer Linux hosts and Cisco switches, this wouldn't be a problem, as there are many excellent and inexpensive iPad Bluetooth keyboards and keyboard/case combos if typing directly on the iPad's screen becomes too tiresome.

Could I get by during crunch time using my finger to navigate the Windows GUI remotely? Of course. Would I want to? Probably not. Will this change in the future? Absolutely. As Windows moves to be more and more PowerShell oriented (along with Microsoft's push to Metro), eventually I expect to see fewer and fewer of these legacy server applications that require you to use Remote Desktop to administer them.

iPad to the Rescue

Outside of that one major negative, my experience using the iPad for "real IT work" is probably best shown by example. Several months ago, I was in the midst of troubleshooting a problem with a virtual machine (VM) host. The host would simply drop dead at random times after approximately one week of running. By "drop dead," I mean there was response to ping, no ability to connect with a kernel debugger over the serial port, no STOP errors, and no kernel dumps. The only way to get the host to come back to life was to power cycle it. The only way to power cycle it remotely was to use the Integrated Lights-Out (iLO) functionality built in to the server. This went on for about three weeks until the root cause was uncovered. (It was a BIOS setting.) When various fixes were being tried during that time, it wasn't uncommon to have to use iLO to access the server at odd hours. One of those odd hours was at 1:00 A.M. while I was walking back to my room at a Holiday Inn during a technology conference.

When my monitoring system told me the system had keeled over, I pulled out the iPad, connected in with VPN, launched Safari, accessed the iLO interface, and restarted the server. This was while I was walking down the hallway. I had also used the iPad heavily all day, and I still had 30 percent of battery life left.

Today I still travel with my laptop. I like having that safety net, especially if tasked with something requiring precision Remote Desktop use. However, I usually reach for the iPad first. Sometime in the next few years, I can see myself not taking the laptop with me.

Despite my laptop potentially going the way of the dinosaur, I can't see giving up my desktop computers. In fact, I find myself using them more. There are many things I just don't want to do without a large monitor, a full-sized keyboard, and a mouse! Their utility can't be beat. This is why many IT pros are experiencing so much trepidation with Windows 8 and Metro on the desktop. I always have dozens of windows open on my desktop machines, flipping through them like a deck of cards.

Give It a Try

If you have an iPad (and considering the large number of them I saw in use at Microsoft TechEd this June, I expect that many of you do), I encourage you to try using it for some of your "real work." Although you might not be able to use it for everything, you'll probably find more uses for it than you originally imagined. In fact, I typed the first draft of this review on my iPad in Evernote. ■

InstantDoc ID 143292

iPad

PROS: Simple to use; huge range of applications and accessories; long battery life; lightweight

CONS: Can't yet completely replace an IT pro's laptop or desktop

RATING: ★★★★★

PRICE: Starting at \$499

RECOMMENDATION: The iPad can't be your only tool, but it's a worthy addition to your arsenal.

CONTACT: [Apple](#) •
800-676-2775 or
408-996-1010

Insights from the Industry



Jeff James

is a former industry news analyst for *Windows IT Pro*. He also was editor-in-chief of *TechNet Magazine* and was an editorial director at the LEGO Company. Jeff has more than 15 years of experience as a technology writer and journalist.

Email



Twitter



Every Major U.S. Company Already Hacked by Chinese Government

Former counterterrorism official Richard Clarke, who once famously apologized to the families of the 9/11 victims by saying, “Your government failed you, those entrusted with protecting you failed you, and I failed you,” is once again back in the headlines, this time for pointing out that America (and its major corporations) are woefully unprepared to fend off cyberattacks from terrorist groups and hostile nation-states such as China.

Clarke was recently [interviewed by Ron Rosenbaum for *Smithsonian Magazine*](#), primarily for an article that focuses on who Clarke believes was behind the Stuxnet cyberattack against Iran in late 2010. Clarke—like many other security experts—points the finger squarely at the United States, hinting that America might have received some assistance from Israeli intelligence services. Here’s a key quote from Clarke about Stuxnet from Rosenbaum’s article:

I think it’s pretty clear that the United States government did the Stuxnet attack . . . I think there was some minor Israeli role in it. Israel might have provided a test bed, for example. But I think that the U.S. government did the attack and I think that the attack proved what I was saying in the book [which

came out before the attack was known], which is that you can cause real devices—real hardware in the world, in real space, not cyberspace—to blow up.

Microsoft Technical Fellow Mark Russinovich writes about terrorists using cyberwarfare to destroy physical machinery in his [fictional novel *Zero Day*](#), but it's clear that Stuxnet—and [possible successors such as Duqu](#)—have turned fiction into reality. Several security experts have suggested that China was [behind the cyberattack on security vendor RSA](#) that netted information about RSA SecurID tokens, information which was then allegedly used to launch attacks against major defense contractors such as Lockheed Martin and Northrop Grumman.

In addition to attacks against US government agencies and defense contractors, Clarke believes that China has invested billions in an attempt to use cyberwarfare to steal secrets from US companies, a strategy that Clarke says the US government won't emulate. Clarke suggests that China's cyber-espionage efforts pose a significant long-term risk to US interests, mainly because China is leveraging cyberwarfare to steal trade secrets from US companies. Clarke told Rosenbaum that his "greatest fear is that, rather than having a cyber-Pearl Harbor event, we will instead have this death of a thousand cuts. Where we lose our competitiveness by having all of our research and development stolen by the Chinese."

Rosenbaum's article presents a chilling outlook for US cybersecurity efforts and should be a must-read for anyone involved in IT security.

—Jeff James

InstantDoc ID 142689

Secure Your Smartphone with the Lookout Mobile Security App

According to a study sponsored by SecurEnvoy, 66 percent of respondents said they suffered from nomophobia—the fear of being without



Blair Greenwood

is an assistant editor for *Windows IT Pro*, *SQL Server Pro*, and *Dev Pro*. She holds a bachelor's degree in technical journalism from Colorado State University. She's worked with Java, C, and Visual Basic programming languages.

Email



Twitter



a mobile device. The study, which included 1,000 respondents from the United Kingdom, also noted that 46 percent took no security measures to protect their device. To learn more about the study, see B. K. Winstead's article "[Smartphone Security & Nomophobia](#)."

It's ironic that this percentage of respondents fails to secure their devices when they fear for being without their smartphone, but I find this statistic to be unsurprising. Although I consider security to be important, I admit that I don't take any measures to secure my own smartphone. I've found that trying to secure my device is too much of a hassle, often-times thinking, "Isn't there an easy way to secure my device?" I assume this is a similar sentiment for most smartphone users.

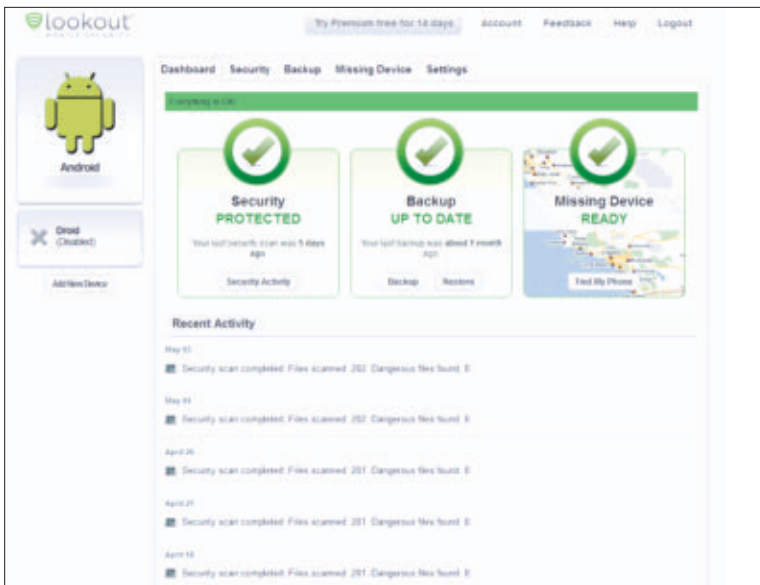
With that said, I discovered a smartphone app that helps secure my device without any additional effort. [Lookout Mobile Security](#) is a free smartphone app that's available for Android and iOS devices. I downloaded the app from the Play Store, and installation was a breeze. I was required to provide my email address to register for a free account with the service and was asked to approve permissions to secure my device through weekly scans and backups.

The free edition of Lookout Mobile Security provides several useful features. Users can choose to have their data backed up daily or weekly. In addition, the app has the ability to automatically scan every app that's installed and perform a full scan each week to protect your device from viruses, malware, and spyware. Contacts are backed up and can be restored to an existing device through Lookout Mobile Security's online dashboard. If you lose your device, you can easily use the Missing Device feature to find your device through the online dashboard. Finally, if you're like me and happen to misplace your smartphone throughout the house—countertops, underneath a pillow or blanket, buried inside your coat pocket—then you can sound an alarm for your device (regardless of whether your device is set on silent mode) through the online dashboard with the Scream feature.

Although the free edition can help locate your device if it has been stolen, this edition isn't particularly useful for protecting your information

in the event of theft. Instead, the premium edition provides much more security from theft and includes the ability to remotely wipe and lock your device. In addition, premium users can back up photos and call history, and data can be restored to a new device. The premium edition takes security a step further and also blocks phishing attempts and protects against malicious websites. The premium edition of Lookout Mobile Security is available for \$2.99 per month or \$29.99 a year.

Overall, I was very pleased with the security that I received with the free edition of Lookout Mobile Security. If you haven't been securing your phone by any means, then take a minute and download this app for the sake of your privacy. See the image gallery to get an idea of how Lookout Mobile Security's features work with its online dashboard.



Slideshow

Lookout Mobile
Security features

Do you secure your smartphone with Lookout Mobile Security or a similar app? I'd love to hear your thoughts on smartphone security. Send me your thoughts on Twitter: [@blair_greenwood](https://twitter.com/blair_greenwood). ■

—Blair Greenwood

InstantDoc ID 143042



**Jason
Bovberg**

Email




Twitter



May the Force Be With Your Car!

Product of the Month

A popular type of vehicle window sticker these days is the family decal that provides a representation of your family via simple illustrations. You've seen them. We were starting to get a little tired of them until someone finally came up with the idea of printing parodies that are gleefully violent or dirty or otherwise hilarious. We can't print some of those here, but you can easily find them online. Our favorite family-friendly versions of these stickers can be found at ThinkGeek: [Star Wars family car decals](#)! These are indispensable for any IT pro's family car.



Send us your funny screenshots, oddball product news, and hilarious end-user stories. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube.



Submit

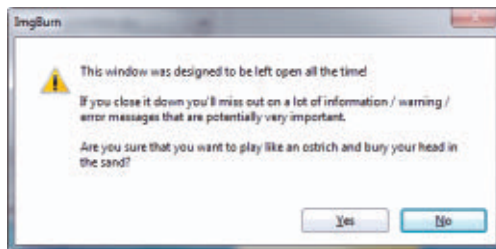


Figure 1: Yes. I am sure of that!



Figure 2: Never need AC power!

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowssitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowssitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowssitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- [Cloud & Virtualization UPDATE](#)
- [Dev Pro UPDATE](#)
- [Exchange & Outlook UPDATE](#)
- [Security UPDATE](#)
- [SharePoint Pro UPDATE](#)
- [SQL Server Pro UPDATE](#)
- [Windows IT Pro UPDATE](#)
- [WinInfo Daily UPDATE](#)

RELATED PRODUCTS

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
www.windowssitpro.com/go/vipsub

SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

1&1 Internet	1
Cisco	10-12
SharePoint Virtual Conference	2, 43
SolarWinds	82-84
WinConnections Fall 2012 Event	6
Windows IT Pro Events Calendar	59
X-IO Technologies	28-30

Vendor Directory

Amazon	109
Android	21, 109
Apple	21, 109, 129
ATEN Technology	107
Attunity	111
Aventura	48
AvePoint	103
Axceler	118
Cisco	115
Citrix Systems	48, 113
Colligo Networks	114
Exclaimer	112

F5 Networks	106
HP	115
K2	114
Lookout	135
Network Instruments	108
Nokia	15, 112
Quest Software	55, 113
Red Hat	55
RES Software	48
Riverbed Technology	111
RSA	135
SolarWinds	105, 115
StillSecure	105
Stratus Technologies	121
Symantec	110
Symplified	107
Telerik	114
ThinkGeek	138
TITUS	113
Veeam Software	55, 106
VMware	8, 52, 60
X-IO Technologies	111